

Three Dimensional Simulations of Space Charge Dominated Heavy Ion Beams with Applications to Inertial Fusion Energy.

By

David Peter Grote
BS, University of Dayton, 1987
MA, University of California at Davis, 1989

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Applied Science

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Committee in Charge

1996

Acknowledgements

I would like to first thank my wife, Cecilia, for getting me through these last few years. I want also to thank Pat Livingston for nagging me to write, and for every time she said “Get that sucker writ”. I would like to thank my research supervisor, Dr. Alex Friedman, for all of his help and enthusiasm. I would also like to acknowledge the help of Drs. Bruce Langdon, Irving Haber, Simon Yu and others of the HIF programs at LLNL and LBL.

David Peter Grote

February 1994

Applied Science

Three Dimensional Simulations of Space Charge Dominated Heavy Ion Beams with
Applications to Inertial Fusion Energy.

Abstract

Heavy ion fusion requires injection, transport and acceleration of high current beams. Detailed simulation of such beams requires fully self-consistent space charge fields and three dimensions. Our code, WARP3D, has been developed for this purpose. It is a particle-in-cell plasma simulation code optimized to work within the framework of an accelerator's "lattice" of accelerating, focusing, and bending elements. The code has been used to study several test problems and for simulations and design of experiments. Two applications presented here are drift compression experiments performed on the MBE-4 facility at LBL and design of the electro-static quadrupole injector for the proposed ILSE facility.

With aggressive drift compression on MBE-4, anomalous emittance growth was observed. Simulations were carried out to examine possible causes. The conclusions of those simulations were that essentially all of the emittance growth is the result of external forces on the beam and is not the result of internal beam space-charge fields. The dominant external forces are the dodecapole component of focusing fields, the image forces on the surrounding pipe and conductors, and the octopole fields that result from the structure of the quadrupole focusing elements.

The goal of the design of the electro-static quadrupole injector is to produce a beam of as low emittance as possible. The simulations show that the dominant effects that increase the emittance are the non-linear octopole fields and the energy effect (fields in the axial direction that are off-axis). Injectors were designed that

minimized the beam envelope in order to reduce the effect of the non-linear fields. Some alterations to the quadrupole structure that reduce the non-linear fields further were examined. Finally, comparisons were done with a scaled experiment. These resulted in very good agreement.

Contents

Acknowledgements	ii
List of Tables	6
List of Figures	7
1 Introduction	1
1.1 Conventions	4
1.2 Emittance	4
1.3 Focusing Fields	5
I The Code	7
2 Methods Used in WARP3d	8
2.1 Particles	10
2.1.1 Residence Corrections	10
2.1.2 Particle Loader	16
2.1.3 Finite or Cigar Shaped Beam	21
2.1.4 Calculation of Particle Moments	27
2.2 Fields	30
2.2.1 Beam Self-Fields	30
2.2.2 External Fields	30
2.2.3 Lattice Description	31

2.2.4	Internal Lattice Description	32
2.2.5	Field Solutions with Internal Conductors	33
2.3	Optimizations, Plot Handling, and the Basis System	41
2.3.1	Techniques Used for Optimization	41
2.3.2	Plot Handling in WARP	42
2.3.3	Basis	43
II	Applications	44
3	Long Time Behavior of Finite Length Beams	45
4	Effects of Round and Square Pipes	51
5	Drift Compression on MBE-4	54
5.1	MBE-4 Results	54
5.2	WARP3d Model of MBE-4	58
5.3	WARP3d Results	62
5.3.1	Dodecapole and Image Fields	71
5.3.2	Inclusion of Internal Conductors	76
5.4	Conclusions	82
6	Electrostatic Quadrupole Injector (ESQ)	87
6.1	Injector Description	87
6.2	Description of Computer Runs	89
6.3	ILSE Scale Injector	92
6.3.1	1 MV Diodes	99
6.4	Self-Consistent Corrections and Optimization	100
6.5	Comparison With Experiment	105
6.5.1	500 kV Diode	105
6.5.2	1 MV Diode	107
6.6	Conclusions of ESQ Injector Study	109

7	Summary	114
A	Envelope Calculation	116
B	Behavior of a Beam in Pure Non-Linear Fields	119
	Bibliography	125

List of Tables

2.1	Examples of binary or bit reversed numbers	18
2.2	Interaction of quadrupoles with capacity matrix field solver.	35
3.1	Parameter list of runs used to study long time behavior of beams. . .	46
5.1	Results of the MBE-4 drift compression experiment.	55
6.1	Parameters for ILSE injector and scaled injector experiment.	93
6.2	Quadrupole voltages for .5 MV scaled injector.	103
6.3	Results of the .5 MV scaled injector with varying quadrupole voltages.	106
6.4	Results of the .5 MV scaled injector with varying initial beam radius.	106
6.5	Initial data for the .5 MV scaled injector.	107
6.6	Initial data for the .5 MV scaled injector.	109

List of Figures

2.1	Diagram of a particle entering a gap.	13
2.2	Cigar velocity distribution function.	25
2.3	Cigar beam	28
2.4	Comparison of axial electric and ear fields	29
2.5	Focusing fields versus number of conductor points.	36
2.6	Comparison of the envelope and emittance of an ESQ injector with and without sub-gridscale boundary conditions.	41
3.1	Uniformly focused cigar beam after 120 m.	47
3.2	Initial and final transverse profiles of an axially confined beam.	48
3.3	Emittance time histories of an axially confined beam	49
3.4	Initial and final transverse profiles of an emittance dominated, axially confined beam.	50
3.5	Emittance histories for an emittance dominated beam.	50
4.1	Comparison of beams in round and square pipes.	53
5.1	Experimental results for the mildly compressed beam.	56
5.2	Experimental results of the aggressively compressed beam.	57
5.3	Comparison of field components with anti-symmetric and nonanti- symmetric quadrupole voltages.	60
5.4	Time history of the line charge at mid-pulse of an ideal compressed beam.	63
5.5	Initial distribution before compression.	64
5.6	Final distribution after compression with 5% velocity tilt.	65

5.7	Final distribution after compression with 10% velocity tilt.	66
5.8	Final distribution after compression with 15% velocity tilt.	67
5.9	Time history of ϵ_x and $2X_{\text{rms}}$ at the center of an ideal compressed beam.	68
5.10	$x - x'$ space of aggressively compressed ideal beam.	69
5.11	Self-potential of a beam in MBE-4 on axis.	70
5.12	Time history of the ϵ_x and $2X_{\text{rms}}$ of an aggressively compressed beam with an applied dodecapole field.	72
5.13	$x - x'$ space of an aggressively compressed beam with an applied do- decapole field.	73
5.14	Transverse profile of the beam undergoing aggressive compression.	74
5.15	Line charge at mid-pulse of a compressing beam in a small pipe.	75
5.16	Comparison of the self-potential of a beam close and far from the wall.	76
5.17	$x - x'$ space and mid-pulse ϵ_x time history of a beam in a small pipe.	77
5.18	$x - x'$ space and mid-pulse ϵ_x time history of a beam in a small pipe with dodecapole.	78
5.19	Mid-pulse ϵ_x time histories with images on quadrupoles.	79
5.20	Dodecapole field with quadrupole rods.	80
5.21	ϵ_x and $2X_{\text{rms}}$ at mid-pulse of a beam in self-consistent quadrupole rods.	80
5.22	Mid-pulse line charge time history of an aggressively compressed beam in the inter-digital structure.	81
5.23	ϵ_x and $2X_{\text{rms}}$ at mid-pulse of a beam in the inter-digital structure.	82
5.24	$x - x'$ trace space at sixteen place along the beam.	83
5.25	$y - y'$ trace space at sixteen place along the beam.	84
5.26	Beam x profile and trace space with symbols.	85
5.27	Beam y profile and trace space with symbols.	85
6.1	Diagram of an inter-digital quadrupole.	88
6.2	The field grid overlaying part of the system.	91
6.3	Initial ILSE injector results.	94
6.4	ILSE injector phase space.	95

6.5	Injector results with V_{42} cancelled.	96
6.6	Fourth order fields in the ILSE injector.	96
6.7	Axial phase space in ILSE injector.	97
6.8	Injector results with energy effect correction.	99
6.9	Injector results with both corrections.	100
6.10	Injector result with a 1 MV diode.	101
6.11	Scaled .5 MV injector with optimized voltages.	102
6.12	Scaled .5 MV injector with optimized voltages and rod offsets.	104
6.13	Comparison of multipole fields in the scaled injector with and without offsets.	105
6.14	Comparison of WARP3d and experiment for the 500 kV scaled injector.	108
6.15	Comparison of WARP3d and experiment for the 1 MV scaled injector.	110
6.16	Comparison of experiment and simulation of 30 kV diode injector. . .	111
6.17	Comparison of experiment and simulation of 30 kV diode injector. . .	111
6.18	Comparison of experiment and simulation of 15 kV diode injector. . .	112
6.19	Comparison of experiment and simulation of 15 kV diode injector. . .	112
B.1	Effect of V_{42} on a cold beam.	120
B.2	Effect of V_{44} on a cold beam.	122
B.3	Effect of V_{40} on a cold beam.	123
B.4	Effect of V_{66} on a cold beam.	124

Chapter 1

Introduction

Creating fusion in the laboratory and ultimately in a power generating plant has long been a goal of scientists worldwide. In the inertial confinement fusion, ICF, approach, a pellet of fusionable fuels is compressed to a high density by impinging laser beams. If the density and the temperature are high enough, the fuel will begin to fuse, releasing energy. If the density is maintained, the fuel will be heated enough to be completely burned. In 1974, Maschke proposed beams of high energy heavy ions as an alternative driver for ICF[1]. That led to the initiation of the heavy ion fusion (HIF) program.

The use of particle beams to compress the fusible targets offers advantages. Much of the technology needed to produce the particle beams has already been developed. Experimental and theoretical scientists have gained much experience in building large scale particle accelerators for use in high-energy particle physics. Particle accelerators have long lifetimes. In a reasonable lifetime of a power generator, the driver will have to produce on the order of 10^{10} shots. Particle accelerators also have high repetition rates and an HIF driver would have no difficulty in producing the required several shots per second. Among other advantages of particle beams are the final focusing lenses. These focus the beam down onto the target in the reaction chamber. Lasers would require a physical lens or a mirror to effect the final focus. This would be subject to degradation from the debris produced in the fusion reaction. Particle

beams are focused by fields which are produced by conductors outside of the line of sight from the entrance to the reaction chamber so there would be no physical lens that could be degraded.

This leaves the choice of the ion mass. With indirectly driven targets, the beam strikes a radiator which then gives off X-rays which then strike the fuel pellet itself. When the particle beam strikes the target, it must react with the radiator only. As much of the beam energy as possible needs to be deposited in the radiator, and the X-rays need to be able to escape from it. Thus, penetration depth of the beam ions into the radiator is critical. Light ions have a larger penetration depth at a given energy than heavy ions. In order to prevent the light ions from heating too much matter and so wasting energy, they need to have fairly low energy. But then, to obtain enough energy to compress the target, very high beam current densities would be required. Heavy ions can have much higher energies, several to ten GeV, and not penetrate too far. The currents needed are much lower, on the order of kiloAmperes. These lower currents are much easier to produce and control. Also, with light ions, the large amount of charge deposited on the target will produce significant electric fields, leading to possible difficulties.

The typical scenario for a heavy ion fusion driver begins with a number of injectors, on the order of 10 to 100, that produce long, on the order of 20 μ s pulse length, beams with an energy of several MeV and a current of about an Ampere. These beams are combined into a smaller number of beams and accelerated in either a linac or a circular accelerator, both using induction accelerating modules. To produce a higher line charge density, the beams are compressed axially. This is accomplished by imparting to the beam a head-to-tail energy, or velocity, gradient. The beam is thereby shortened, producing a higher line charge density. The beam is compressed down to nearly a 10 nanosecond pulse length and up to several thousand Amperes. The beam is then bent to go into the reaction chamber and sent through the final focus. All in all, the beam will be transported over at least several kilometers.

The target is of order a few centimeters in diameter. In order for it to be possible

to focus the beam down onto the target, a low transverse beam temperature, or emittance, must be maintained. (Emittance is the measure of phase space area and is described in more detail shortly.) This represents the major difficulty in the design of a driver. The various manipulations of the beam and the long distance over which the beam is transported all tend to increase the transverse temperature. Accurate and detailed analysis is necessary in the design of a driver in order to minimize the growth of emittance.

The goal of this thesis research was to develop a three dimensional code to simulate the transport and acceleration of high current beams. A three-dimensional code was needed since many of the beam manipulations are inherently three-dimensional. The injection requires transverse focusing (confinement) of the beam, possibly with alternating gradient (AG), or quadrupole, lenses and acceleration of the beam. Beam merging involves beams that do not travel parallel to the machine centerline. Longitudinal compression involves the axial gradients and the expansion of the beam transversely within an AG focusing system. The interaction of the transverse and longitudinal physics is important in the transport through an AG focusing system.

In order to simulate high current beams correctly, the beam self-fields must be included. In a driver scale accelerator, the self-fields from the space-charge dominate over the beam temperature; nearly all of the focusing is counteracted by the electrostatic repulsion of the beam particles. The particle-in-cell method (commonly used to simulate plasmas) was chosen. A small number of particles (compared to the number of real particles) are used to represent the beam and the self-fields are calculated self-consistently from them. This method was chosen in part since with it all of the beam physics can be calculated from first principles. It is also faster than other particle methods, such as the particle-particle method. Furthermore, with its extensive use in plasma simulations, the method has become very well understood.

The code that was developed was used in several applications: long time transport of a beam in an alternating gradient focusing channel, comparison of two types of beam pipes and their image charges, simulation of drift compression on MBE-4[2], and

simulation and design of the electrostatic quadrupole injector for the induction linac systems experiment (ILSE)[3]. These applications both validated the code against experiment and demonstrated the unique capabilities and the utility of the code.

1.1 Conventions

In a straight system, the coordinates are a right handed Cartesian system with x on the horizontal, y on the vertical, and z along the centerline of the beam. The beam is always assumed to move in the positive z direction. Derivatives with respect to time are for some purposes converted into derivatives with respect to z by division of the average axial velocity of the beam. Derivatives with respect to z are denoted by primes, and with respect to time are denoted by dots. For example,

$$x' = \frac{dx}{dz} = \frac{dx}{dt} \frac{dt}{dz} = \frac{\dot{x}}{v_z} \quad (1.1)$$

In any finite differenced equations, superscripts refer to time steps and subscripts refer to grid cell index, except as noted.

1.2 Emittance

Emittance, ϵ , is the measure of the area of the beam in “trace” space, or (x, x') space, and is related to the transverse temperature. The normalized emittance, ϵ_N , is defined as the area of the beam in the canonical phase space (x, p_x) and is related to emittance via the beam axial velocity, $\epsilon_N = \gamma\beta\epsilon$. ($\beta = \frac{v_z}{c}$ and $\gamma = \frac{1}{\sqrt{1-\beta^2}}$ are the usual relativistic terms and c is the speed of light in a vacuum.) In the Kapchinski-Vladimirskij (KV)[4] distribution, the only known equilibrium distribution in an alternating gradient accelerator, the beam uniformly fills an ellipse in phase space. The emittance is given by the product of the major and minor radii of that ellipse and π , which is equivalent to the area. For beams that are not uniform or do not form an ellipse, emittance is defined alternatively as a root mean square (RMS) “edge” emittance.

$$\epsilon_{\text{RMS}} = 4\sqrt{\langle(x - \langle x \rangle)^2\rangle\langle(x' - \langle x' \rangle)^2\rangle - \langle(x - \langle x \rangle)(x' - \langle x' \rangle)\rangle^2} \quad (1.2)$$

Here $\langle \rangle$ means an average over the distribution. The units of ϵ_N are commonly “ π -meter-radians” or “ π -mm-mr”, with the “ π ” explicitly included to show that it has not been absorbed into the numerical value. For a KV beam, the RMS expression is equivalent to the phase space area. The RMS emittance is sometimes a pessimistic measure. Phase space may be bent and distorted in a way that can be “undone” so the true phase space area has not changed, but by the RMS definition, the emittance will have increased. If such distortions were undone, the RMS emittance would decrease. True emittance, in general, will only decrease under very unusual circumstances.

The Liouville theorem states that on a microscopic scale the occupied phase space area remains constant. When all of the forces are linear, uniformly filled ellipses map into ellipses, and the emittance remains constant as well. When non-linear forces are present, the emittance can increase or decrease (the system is time-reversible). The possible non-linear forces include (among others) the self-fields of a non-uniform beam, fields from imperfect focusing structures, and image fields on the surrounding conductor.

1.3 Focusing Fields

The focusing fields in an alternating gradient focusing system consist of quadrupole fields that focus the beam in one direction and defocus it in the other. They are generally arranged so that each transverse direction is alternately focused and defocused. This is referred to as the “FODO” lattice, for focusing, drift, defocusing, drift (“O” stands for open, for an opening between conductors, or zero, for zero fields in the drift regions). The unit is repeated. The fields can either be electrostatic or magnetic. The quadrupole potential in the electrostatic case has the form

$$\phi = \phi_0(x^2 - y^2) = \phi_0 r^2 \cos 2\theta, \quad (1.3)$$

where x and y or r and θ are the transverse coordinates.

With imperfect electrostatic focusing elements, non-linear fields may be present.

The potential is written in multipole components as

$$\phi(r, \theta, z) = \sum_{k \geq l \geq 0} \phi_{kl}(z) = \sum_{k \geq l} V_{k,l}(z) \left(\frac{r}{R}\right)^k \cos(l\theta) \quad (1.4)$$

where $V_{k,l}(z)$ are the field strengths, and R is the quadrupole aperture, and $k - l$ is even. The quadrupole field is given by $\phi_{2,2}(z)$.

The requirements on k and l can be seen by putting the summation back into the Laplace equation in cylindrical coordinates.

$$\nabla^2 \phi = \frac{1}{r} \frac{\partial}{\partial r} r \frac{\partial}{\partial r} \phi + \frac{1}{r^2} \frac{\partial^2}{\partial \theta^2} \phi + \frac{\partial^2}{\partial z^2} \phi = 0 \quad (1.5)$$

After grouping like terms in r and θ , the resulting expression is

$$\begin{aligned} & \sum_{l \geq 0} r^{-2} \cos(l\theta) (-l^2) V_{0,l} + \\ & \sum_{l \geq 0} r^{-1} \cos(l\theta) (1 - l^2) V_{1,l} + \\ & \sum_{l \geq 0} \sum_{k \geq 2} r^{k-2} \cos(l\theta) [(k^2 - l^2) V_{k,l} + V''_{k-2,l}] = 0. \end{aligned} \quad (1.6)$$

To keep $\nabla^2 \phi$ well behaved at $r = 0$, $V_{0,l}$ must vanish for all $l > 0$ and $V_{1,l}$ must vanish for all l except $l = 1$. Requiring each term in the sum over k to be zero independently gives a recursion relation for $V_{k,l}$.

$$V_{k,l} = V''_{k-2,l} / (l^2 - k^2) \quad (1.7)$$

So, since $V_{0,l}$ is zero for $l > 0$, $V_{k,l}$, with k even, is zero for all $k < l$. $V_{k,k}$, with k even, and hence $V_{k,l}$ with k even and $k > l$, is not zero since both the numerator and denominator on the right hand side of the recursion relation vanish, leaving $V_{k,k}$ arbitrary. But the $V_{k,l}$ with k even and l odd are zero for $k > l$ since at no place in the recursion does the denominator $l^2 - k^2$ vanish. Starting with the necessity that $V_{1,l}$ must vanish for all l except $l = 1$, the terms with $k < l$ and $k - l$ odd vanish similarly. The remaining nonzero terms all have $k \geq l \geq 0$ and $k - l$ even.

With quadrupole, or two plane, symmetry, neither k nor l will be odd since $\phi(x, y, z) = \phi(-x, y, z)$. When l is odd, $\phi(x, y, z) = -\phi(-x, y, z)$.

Hereafter, the comma in the subscript of $V_{k,l}$ will be removed and it will be written as V_{kl} . The only terms which are significant and are studied have both k and l less than ten so there will be no confusion with the comma removed.

Part I

The Code

Chapter 2

Methods Used in WARP3d

The basis of the WARP code is the electrostatic particle-in-cell (PIC) method[5]. This method uses the Lorentz equation of motion to advance in time macro-particles, which represent many real particles. The only self-field considered is the electrostatic field. This is calculated self-consistently on a three dimensional, rectangular mesh that holds a discrete representation of the charge density and the electrostatic potential. The charge density is calculated from the particles via a linear interpolation of the macro-particles onto the mesh. Three dimensional fast Fourier transforms are used to solve for the electrostatic potential from the charge density using Poisson's equation. The macro-particles are advanced in time using a combination of the "leap frog" and "isochronous leap frog" methods[7]. Each time step goes through the following pattern:

1. The charge of the macro-particles is deposited onto the mesh, using trilinear weighting.
2. The electrostatic potential is calculated.
3. The electric fields are interpolated from the mesh to the macro-particles.
4. The velocity and then the position of the macro-particles are advanced.

Hereafter, the term “particle” should be interpreted to mean macro-particle unless explicitly stated otherwise.

For an accelerator code, additions to the usual PIC method were needed. These include external fields like the focusing, bending and accelerating fields, a lattice description, accelerator-specific diagnostics, and conductors within the field mesh. Various techniques are used in the code for optimization.

One note on the “leap frog” method is the splitting of steps. In the usual method, the velocities are known at half integral time steps while the positions are known at integral time steps. The time discrepancy can lead to errors when diagnostic plots are made or when moments are calculated. When diagnostics are made, typically every five or ten timesteps, the velocity advance is split into two steps and an “isochronous” advance is used. The first half step synchronizes (in time) the velocity with the position. The diagnostics can then be done. The second half step of the velocity advance is then done at the beginning of the next timestep. When no diagnostics are being taken, the code uses the unsplit method for efficiency.

I had my hand in the development of nearly all parts of the code. A number of the sections described in this thesis were developed and implemented entirely by me. (By implemented, I mean I wrote the fortran code.) These include the residence corrections with acceleration, the moments calculations, and the sub-gridscale boundary placement in the field solver. I developed all of the particle loader and injection except the “cigar” beam loader which is based off of a colleague’s work, and the random number generators which are extensions of what was obtained in the literature; I implemented all of it. The capacity matrix and SOR field solver were taken from the literature with several new extensions added by me and written in their entirety by me. In the remaining sections of code described in this thesis I played a significant role.

2.1 Particles

The following sections describe the handling of particles.

2.1.1 Residence Corrections

It is desirable to take as few, and therefore as large, steps as is possible. In general, for a travelling beam in a lattice of hard-edged elements, the limiting factor in the size of the step is getting the correct impulses from the elements. Since the particles all do not step at the same axial positions, some may “land inside” a lattice element on more steps than others, leaving them with different impulses. Small steps would be needed to decrease the errors from this effect. We get around this by performing “residence corrections”, that is, correcting the forces applied from each element by scaling them by the fraction of the timestep the particle spent inside the element. This greatly relieves the timestep size constraint, reducing the number of steps needed within the elements to typically ten or fewer.

Residence correction may be done for only the particle velocity, or for both the position and velocity. For the hard-edged focusing and bending elements, which primarily affect the transverse velocity and position, only the correction to the transverse velocity is required for sufficient accuracy. No correction is needed in the axial velocity and position and the correction in the transverse position is small. The correction scales as the time step size, Δt , squared times the transverse velocity which is small compared to the axial velocity. The acceleration lattice requires a correction on the axial position of the particle as well as a correction to the velocity.

Application of the correction to the velocity only is the easier case. It is done when the velocity would be affected by the element, that is, if the particle would enter or leave the element between the time levels $n - 1/2$ and $n + 1/2$ (the velocity is usually known only at half integral steps). The velocity correction consists of an adjustment to the applied field according to the fraction of the time step that was spent within the element. This fraction is calculated as a ratio of the axial distance

travelled within the element to the axial distance travelled during the time step. The ratio of distances can be used since the axial velocity is nearly a constant during the time step. For a full time step, the axial distance travelled is given by $v_{\text{beam}}\Delta t$ and is centered about the particle's position. This distance represents the range over which the forces are integrated. For the steps during which diagnostics are taken, this distance is given by $\frac{1}{2}v_{\text{beam}}\Delta t$, covering the range appropriate to the part of the advance being done. Define Δt_l (for “left” since our beams usually travel left-to-right) as the change in time from the most recent time that the velocity is known to the current time. This will be either 0 or $-\frac{1}{2}\Delta t$ depending on whether or not the velocity is synchronized with the position. Define Δt_r as the change in time from the current time to the time the velocity will be known. This will be either 0 or $\frac{1}{2}\Delta t$, depending on whether or not the velocity will be synchronized with the current time. The lower end of the range of position is then given by

$$z_l = z_p + v_z \Delta t_l, \quad (2.1)$$

and the upper end is given by

$$z_r = z_p + v_z \Delta t_r, \quad (2.2)$$

where z_p is the current axial position of the particle and v_z is the most recent axial velocity of the particle. When the particle enters the element, the fraction of the step in the element is given by

$$f = \frac{z_r - z_{\text{entrance}}}{v_z(\Delta t_r - \Delta t_l)}, \quad (2.3)$$

and when it leaves, by

$$f = \frac{z_{\text{exit}} - z_l}{v_z(\Delta t_r - \Delta t_l)}, \quad (2.4)$$

where z_{entrance} and z_{exit} are the positions of the entrance and exit of the element respectively. The applied field is then multiplied by this fraction.

Since the velocity lags behind the position by one half a time step, the effect of the element will also lag. It can happen that the particle will enter the element on the step before its velocity is affected. Likewise, a particle can have exited the element

and on the next step have its velocity still affected by it. The total amount that each of the particles are affected must still be the same.

When there is axial acceleration, the residence correction becomes much more complicated. The corrections needed for the position and the end points of the velocity range are more difficult to determine since the axial velocity changes significantly over the course of the step. To calculate the residence corrections for the position, the position at the end of the step is first calculated by integrating the equation of motion over one time step starting with the known current values. The velocity $v^{n+\frac{1}{2}}$ is also calculated from the equation of motion. The calculated position is then compared to what the position would be if advanced with leap frog using the calculated velocity, $v^{n+\frac{1}{2}}$. The difference gives the necessary correction. The calculations are estimations since only the accelerating field is included; all other forces, in particular the space charge forces, are ignored.

On entering the gap, the “exact” position at the end of the next step is

$$z^{n+1} = z^n + \int_0^{\frac{\delta_z}{v^n}} v^n dt + \int_{\frac{\delta_z}{v^n}}^{\Delta t} \left(v^n + a_g \left(t - \frac{\delta_z}{v^n} \right) \right) dt, \quad (2.5)$$

where z and v are the position and velocity, n is the time level, δ_z is the distance between the current position and the entrance of the acceleration gap, $z_{\text{entrance}} - z^n$, and $a_g = E_g \frac{q}{m}$ is the acceleration from the gap. This evaluates to

$$z^{n+1} = z^n + v^n \Delta t + \frac{1}{2} a_g \Delta t^2 \left(1 - \frac{\delta_z}{v^n \Delta t} \right)^2. \quad (2.6)$$

Here, v^n is the same as $v^{n-\frac{1}{2}}$ since no acceleration occurs before the particle’s z coordinate reaches the gap.

The calculation of the velocity depends on when the particle enters the gap. Figure 2.1 makes the following explanation more clear. If the particle enters the gap in the first half of the position advance step, case #1, the velocity $v^{n+\frac{1}{2}}$ will be affected by the gap. The velocity is then

$$v^{n+\frac{1}{2}} = v^{n-\frac{1}{2}} + a_g \Delta t \left(\frac{1}{2} - \frac{\delta_z}{v^{n-\frac{1}{2}} \Delta t} \right). \quad (2.7)$$




Figure 2.1: Diagram of a particle entering a gap. The solid circle are the times when the position is known, empty circles are when the velocities are known, i is the current time level. In 1), the particle is entering the gap in the first half of the position advance step; in 2), the particle is entering in the second half of the position advance step.

In case #2, the velocity is unaffected since the particle enters the gap in the second half of the position advance step, so $v^{n+\frac{1}{2}} = v^{n-\frac{1}{2}}$. Now, if we use equation 2.7 to calculate z^{n+1} with leap frog, we have then for case #1

$$\begin{aligned} z^{n+1} &= z^n + v^{n+\frac{1}{2}} \Delta t, \\ &= z^n + v^{n-\frac{1}{2}} \Delta t + a_g \Delta t^2 \left(\frac{1}{2} - \frac{\delta_z}{v^{n-\frac{1}{2}} \Delta t} \right). \end{aligned} \quad (2.8)$$

The difference of this with equation 2.6 gives the correction

$$z_{\text{corr}} = \frac{1}{2} a_g \left(\frac{\delta_z}{v^{n-\frac{1}{2}}} \right)^2. \quad (2.9)$$

For case #2, when the particle enters the gap in the second half of the step, the next position would be given by

$$z^{n+1} = z^n + v^{n-\frac{1}{2}} \Delta t, \quad (2.10)$$

so the correction would be

$$z_{\text{corr}} = \frac{1}{2} a_g \left(\Delta t - \frac{\delta_z}{v^{n-\frac{1}{2}}} \right)^2. \quad (2.11)$$

When the particle leaves, a similar integration is done:

$$z^{n+1} = z^n + \int_0^{\delta_t} (v^n + a_g t) dt + \int_{\delta_t}^{\Delta t} v_g dt \quad (2.12)$$

where δ_t is time spent in the gap which is assumed to be less than Δt and v_g is the velocity at the time of exit from the gap. The time spent in the gap, δ_t , can be calculated from the distance in the gap.

$$z_{\text{exit}} - z^n = v^n \delta_t + \frac{1}{2} a_g \delta_t^2. \quad (2.13)$$

This quadratic equation has the solution

$$\delta_t = \frac{1}{a_g} \left[\left((v^n)^2 + 2a_g(z_{\text{exit}} - z^n) \right)^{\frac{1}{2}} - v^n \right]. \quad (2.14)$$

The velocity at the exit of the gap is given by $v_g = v^n + a_g \delta_t$. The velocity, v^n , in terms of the known velocity at time level $n - 1/2$, is given by $v^n = v^{n-1/2} + \frac{1}{2} a_g \Delta t$. With all of this, the position at the next time level, equation 2.12, reduces to

$$z^{n+1} = z^n + v^{n-1/2} \Delta t + \frac{1}{2} a_g (\Delta t^2 + 2\delta_t \Delta t - \delta_t^2). \quad (2.15)$$

For case #1, when the particle exits the gap in the first half of the step, the velocity is given by

$$v^{n+1/2} = v^{n-1/2} + a_g \left(\frac{1}{2} \Delta t + \delta_t \right) \quad (2.16)$$

By the same method as before, this gives as the position correction

$$z_{\text{corr}} = -\frac{1}{2} a_g \delta_t^2. \quad (2.17)$$

For case #2, when the particle exits in the second half, the velocity is

$$v^{n+1/2} = v^{n-1/2} + a_g \Delta t, \quad (2.18)$$

so the correction would be

$$z_{\text{corr}} = -\frac{1}{2} a_g (\Delta t - \delta_t)^2. \quad (2.19)$$

In the limit $\delta_t \rightarrow 0$, as $z^n \rightarrow z_{\text{exit}}$, z_{corr} for case #1 goes to zero as expected since the position advance is completely outside of the gap. As $\delta_t \rightarrow \Delta t$, z_{corr} for case #2 goes to zero as expected since the position advance is completely inside the gap.

The calculation of the velocity correction is made difficult with acceleration because the axial velocity changes significantly over the time step. This affects the fraction of the step spent in the element. Also, the fraction of the step in the element must strictly be calculated as a ratio of the time spent in the element to the step length, Δt . For the calculation of the ratio, the velocity at the same time level as the position, v^n , is required. This includes the possibility that the entrance or exit from the gap will be crossed in the half step between $n - \frac{1}{2}$ and n . There are four possible cases. The first is that the half step is completely outside a gap. In this case, v^n is the same as $v^{n-\frac{1}{2}}$. The second case is when the half step is entirely within the gap. The relation here is $v^n = v^{n-\frac{1}{2}} + \frac{1}{2}a_g\Delta t$. The other two cases are when either an entrance or an exit is crossed. Simple derivations give the result

$$v^n = \left[(v^{n-\frac{1}{2}})^2 + 2a_g(z^n - z_{\text{entrance}}) \right]^{\frac{1}{2}} \quad (2.20)$$

on entrance, and

$$v^n = \frac{1}{2} \left(v^{n-\frac{1}{2}} + \frac{1}{2}a_g\Delta t + \left[\left(v^{n-\frac{1}{2}} + \frac{1}{2}a_g\Delta t \right)^2 - 4a_g(z_{\text{exit}} - z^n) \right]^{\frac{1}{2}} \right) \quad (2.21)$$

on exit. The entrance is known to be crossed when it is within the range $z_1 < z_{\text{entrance}} < z^n$. Here, z_1 is given by $z^n - \frac{1}{2}v^{n-\frac{1}{2}}\Delta t - \frac{1}{2}a_g\Delta t^2$. This is the distance that would be travelled if the movement was completely in the gap. The exit is known to be crossed when it is within the range $z_2 < z_{\text{exit}} < z^n$, where z_2 is given by $z^n - \frac{1}{2}v^{n-\frac{1}{2}}\Delta t$. This is the distance travelled completely outside the gap.

Once v^n is known, the calculation of the time of entrance or exit from the gap can be done. This then gives the time spent in the gap and the desired ratio. There are six cases here. If the step is completely outside, of course, the ratio is zero, and if the step is completely inside, the ratio is one. There are two cases for the entrance. One when the entrance is in the first half of the step, the second when the entrance is in the second half of the step. There are the same two cases for the exit. To make the coding simpler, the first two cases are included implicitly by use of max and min functions. That is, if the ratio calculated is less the zero, then it is zero, and if is

calculated to greater than one, it is then one. The remaining four are shown. When $z^n \leq z_{\text{entrance}}$,

$$\text{ratio} = \max \left[\frac{1}{2} + \frac{z^n - z_{\text{entrance}}}{v^n \Delta t}, 0 \right]. \quad (2.22)$$

When $z_{\text{entrance}} < z^n < z_{\text{mid}}$, where z_{mid} is the position of the middle of the gap,

$$\text{ratio} = \min \left[\frac{1}{2} - \frac{v^n}{a_g \Delta t} \left[\left(1 - \frac{2a_g(z^n - z_{\text{entrance}})}{(v^n)^2} \right)^{\frac{1}{2}} - 1 \right], 1 \right]. \quad (2.23)$$

When $z_{\text{mid}} < z^n < z_{\text{exit}}$,

$$\text{ratio} = \min \left[\frac{1}{2} + \frac{v^n}{a_g \Delta t} \left[\left(1 + \frac{2a_g(z_{\text{exit}} - z^n)}{(v^n)^2} \right)^{\frac{1}{2}} - 1 \right], 1 \right]. \quad (2.24)$$

At last, when $z_{\text{exit}} < z^n$,

$$\text{ratio} = \max \left[\frac{1}{2} - \frac{z^n - z_{\text{exit}}}{v^n \Delta t}, 0 \right]. \quad (2.25)$$

Here, $\min[]$ and $\max[]$ have the usual meanings. As with the position correction, all of the square roots can be removed by making the approximation that the acceleration per step is small, $a_g \Delta t / v^{n-\frac{1}{2}} \ll 1$.

For purposes of precision in the calculation in the code, expressions of the form $(a^2 + x)^{\frac{1}{2}} - a$, where $x \ll a$, are rearranged. The numerator and denominator are multiplied by the conjugate $(a^2 + x)^{\frac{1}{2}} + a$. This changes the expression to the form

$$\frac{x}{(a^2 + x)^{\frac{1}{2}} + a}. \quad (2.26)$$

This removes the subtraction of similar numbers which is prone to errors.

The algorithm was verified by using it to track particles in one dimension through a finite length gap and comparing the results with the analytic answer. The results were identical within the precision of the computer.

2.1.2 Particle Loader

In writing the particle loading routine, two issues had to be taken into consideration: what distributions to model, and what kind of random or pseudo-random numbers should be used. The random numbers are discussed first.

It is desirable to load particles in a fashion that minimizes noise in the distribution function. The noise is caused by clumping of particles and by voids, spaces with few particles. Purely random and pseudo-random number generators tend to produce distributions with significant noise. Two non-random “quiet start” number generators were implemented. These generators produce numbers that are nearly uniformly distributed but with a complex ordering.

The first is a higher dimensional generalization of the two-dimensional “Fibonacci” generator. The numbers are chosen using modulus one of a factor times an index. In three dimensions, this looks like

$$\begin{aligned} x_1 &= i/N - 1/2N, \\ x_2 &= \left[\frac{(i-1/2)*g_2}{N} \right] \bmod 1, \\ x_3 &= \left[\frac{(i-1/2)*g_3}{N} \right] \bmod 1, \\ i &= 1, \dots, N, \end{aligned} \tag{2.27}$$

where N is the total number of points and the g ’s are numbers less than N . For two dimensions, the optimum distribution is given when N is the Fibonacci number α_n and g_2 is α_{n-1} . For higher dimensions, there is no clear method for choosing good g numbers. See reference [8] for empirically determined tables. Use of this generator is discouraged as we had problems getting smooth enough distributions, and the lack of a good method for determining g ’s limits the use of the method, especially for large N .

The second is a generalization of the two dimensional bit reversed method. The first dimension is a steadily increasing number as above,

$$\begin{aligned} x_1 &= i/N + 1/2N, \\ i &= 0, \dots, N-1, \end{aligned} \tag{2.28}$$

and the rest are given by converting i to different bases, reversing the digits and converting back to decimal. In two dimensions, the second dimension is found by converting to binary, reversing the bits, and converting back to decimal. See table 2.1 for examples. A third dimension would be obtained by converting to base three,

i	binary	reversed	decimal
0	0	.0	.0
1	1	.1	.5
2	10	.01	.25
3	11	.11	.75
4	100	.001	.125
5	101	.101	.625
6	110	.011	.375
7	111	.111	.875
8	1000	.0001	.0625

Table 2.1: Examples of binary or bit reversed numbers

reversing the digits and converting back to decimal. A constant $1/2N$ is then added to keep the numbers nonzero.

Pseudo-random numbers are also used, especially in the thermal velocity distributions. A fourth non-random option can be used as well, whereby the particles are loaded onto a rectangular grid with possible small random adjustments off the grid in the axial direction.

Our experience has shown that the digit-reversed method generally gives the best distributions. The advantage of uniform distributions is lost, though, in longer runs, as the particles mix and the fluctuations rise to a thermal level. See reference [9] for further information on non-random number generators.

Beam Distributions

There are two different longitudinal distributions available, an infinite beam and a finite beam. The axial position and velocity are loaded the same way for both, except that for the finite beam, they are usually transformed to give a rounded end. This transformation is done by the routine which is described in detail in section 2.1.3. The transformation makes the falloff in line charge at the ends of the beam parabolic. The routine produces an array containing the square root of the line charge that is used to scale the transverse dimensions of the beam. The scaling is done in a way

that keeps the tune depression the same as in the rest of the beam.

Except for the grid loading (the fourth option described above), the axial positions steadily increase from the tail of the beam to the head. They are given by

$$z(i) = z_{\min} + Z_{\text{len}} \frac{(i - .5)}{n}, \quad (2.29)$$

where i varies from one to n , and n is the total number of particles. The axial thermal velocity distribution is Gaussian, except when a finite beam is loaded with a rounded end. That distribution is described in section 2.1.3. Once the longitudinal positions are obtained, the envelope for the transverse distributions can be calculated. The loader can either use the output from the envelope calculation, or can load the beam into a cylinder. Appendix A details the envelope calculation. When the axial velocity is loaded, an optional head to tail velocity gradient, or “tilt”, can be added. This is used to initiate axial compression of the beam.

Semi-Gaussian Distribution

Two transverse distributions have been used, a semi-Gaussian distribution that is Gaussian in velocity and has a “top hat” distribution in position, and the Kapchinskij-Vladirmiskij (K-V) distribution [4].

For the semi-Gaussian distribution, the positions and velocities are loaded independently of each other. Different random number generators can be used. The positions are first loaded into a normalized rectangular block. The block is then carved and scaled to the envelope shape as described below. For the cylindrical beam, all of the points outside the cylinder are thrown away. Carving the beam to fit the calculated envelope can be done in one of two ways. The first is the streamline method, whereby the block is carved into a cylinder and then scaled to fit the envelope. The other is the stripe method; the block is carved directly to the shape of the envelope. The streamline method is the preferred method since the axial variation in density more closely matches that of a real beam and the line charge density is constant. With the stripe method, the density is uniform and the line charge is not constant, but depends on the area of the beam in the transverse plane.

The thermal part of the velocities are Gaussian. The coherent transverse velocities associated with motion in the alternating-gradient fields, which are scaled to the envelope and normalized positions, are then added.

Hollow beam distributions can be used. The “top hat” distribution is transformed it into the distribution

$$f(r_2) = \begin{cases} f_0(\frac{r}{r_{max}})^2, & \text{if } r < r_{max}/2; \\ f_0(1 - (\frac{r}{r_{max}})^2)/3, & \text{if } r > r_{max}/2. \end{cases} \quad (2.30)$$

where f_0 is the normalization constant. The thermal velocity is not affected by the transformation.

Kapchinskij-Vladirmiskij Distribution

With the K-V distribution, the transverse positions and thermal velocities are loaded onto a four dimensional ellipsoid, represented by the normalized delta function $\delta(x^2 + y^2 + x'^2 + y'^2 - 1)$. Only three random numbers are generated, the fourth is calculated using the delta function. The positions and velocities are each loaded into ellipses to give uniform distributions, the two dimensional projections of the 4-D ellipsoid. The first random number is used to calculate the radius in the ellipse. The second is the angle for the position, and the third is the angle for the velocity. The alternating gradient velocities and positions, which are scaled to the envelope, are then added.

After the beam is loaded, the particle number and weights can be set. The charge and mass of the real particles represented by the code particles are stored. The conversion factor between real particles and code particles is also calculated. It is the ratio of the number of real particles to the number of code particles. The factor has to be adjusted for the finite beam since the middle of the beam is slightly compressed from the redistribution during the “cigar” transformation (discussed below). Particles are moved out of the ends of the beam to give a lower line charge density. The conversion factor is mostly used in the deposition of the particles to the charge density grid to give the charge density in MKS units.

Injection

Injection is handled by routines that are separate from the particle loading routines just described. The injection implemented is not space-charge-limited. At each time step, a preset number of particles is injected from the plane of injection. The particles are distributed evenly over the axial distance travelled in one time step. That distance includes acceleration by any potential gradients along the beam axis. The particles are distributed as if they were injected uniformly in time and accelerated along a uniform electric field given by the ratio of the potential change across the region of injection to the length of the region.

Currently, the transverse particle distribution is limited to a Semi-Gaussian distribution, uniform in space over an ellipse, and Gaussian in thermal velocity. The positions are determined with the digit-reversed non-random number generator.

The handling of particles in memory takes advantage of the dependence of the memory location on the axial position. Particles are loaded in increasing order in axial position so that particles with larger axial position are located higher in memory. Since beams have positive velocity, the particles which will exit the grid first will be the particles with larger axial position, and therefore higher memory location. The initially loaded particles are moved to the end of the arrays, leaving the beginning of the arrays empty. The particles which are then injected are placed in the arrays just before the already loaded particles, slowly filling up the first part of the arrays. Also, any lost particles are cleared out of the end of the arrays. When the beginning of the arrays are filled, the active particles are moved to the end of the arrays and the process repeats.

2.1.3 Finite or Cigar Shaped Beam

With a finite length beam, the axial positions and velocities can be transformed so that the beam has tapered ends in which the line charge falls off parabolically. The square root of the normalized line charge density is also calculated and is used to scale the transverse positions and velocities. This beam profile resembles a cigar, hence

the name.

The transformation starts with normalized positions and velocities with uniform distributions. They are scaled to have the tapered distribution. The tapered, or “Neuffer”, distribution used is derived in reference [10]. A parabolic falloff in line charge is desired so the axial electric field, which is approximately proportional to $\frac{d\lambda}{dz}$, is linear. The ends of the beam can then be axially confined with externally applied linear fields.

In a long beam, the line charge is parabolic in the ends and constant in the center. When the length of the whole beam is 1, and the length of each end is z_e , the line charge is given by

$$\lambda(z) = \begin{cases} \lambda_o(1 - \frac{(z_e - z)^2}{z_e^2}), & \text{for } 0 < z < z_e; \\ \lambda_o, & \text{for } z_e \leq z \leq 1 - z_e; \\ \lambda_o(1 - \frac{(z - 1 + z_e)^2}{z_e^2}), & \text{for } 1 - z_e < z < 1. \end{cases} \quad (2.31)$$

The transformation to a cigar shaped beam is accomplished by integrating the line charge density of the cigar distribution and comparing that to the integration of the line charge density of a uniform distribution.

In the range $0 < z < z_e$, first set λ_o to $\frac{3}{2z_e}$ so that the integral of λ from 0 to z_e is 1. Then define the uniform line charge,

$$\lambda_u = \begin{cases} 0, & \text{for } 0 < z < z_e/3; \\ \frac{3}{2z_e}, & \text{for } z_e/3 < z < z_e. \end{cases} \quad (2.32)$$

so it too has a integral over $0 < z < z_e$ equal to 1. Then integrate each and set them equal to each other.

$$\int_0^z \lambda dz = \int_{\frac{z_e}{3}}^{z_u} \lambda_u dz \quad (2.33)$$

Here, z_u is a location in the uniform distribution and z is the location to which it is transformed. Expanding this out and rearranging gives a cubic equation in z .

$$\frac{(z - z_e)^3}{z_e^3} - 3\frac{z - z_e}{z_e} + 3\frac{z_u - z_e}{z_e} = 0 \quad (2.34)$$

The roots were found in a standard reference, and are all real. Only one falls within the desired range.

$$\begin{aligned}\alpha &= \frac{1}{3} \tan^{-1} \left[\sqrt{1 - \left(\frac{3}{2} \frac{z_u - z_e}{z_e} \right)^2} / \left(\frac{3}{2} \frac{z_u - z_e}{z_e} \right) \right] \\ \beta &= -\cos \alpha + \sqrt{3} \sin \alpha \\ z &= (1 + \beta)z_e\end{aligned}\tag{2.35}$$

The transformation is done similarly in the other end of the beam, in the range $1 - z_e < z < 1$. The uniform line charge is defined as

$$\lambda_u = \begin{cases} \frac{3}{2z_e}, & \text{for } 1 - z_e < z < 1 - z_e/3; \\ 0, & \text{for } 1 - z_e/3 < z < 1. \end{cases}\tag{2.36}$$

Equating the integrals as above,

$$\int_{1-z_e}^z \lambda dz = \int_{\frac{z_e}{3}}^{z_u} \lambda_u dz,\tag{2.37}$$

gives a cubic equation in z ;

$$\frac{(z - (1 - z_e))^3}{z_e^3} - 3 \frac{z - (1 - z_e)}{z_e} + 3 \frac{z_u - (1 - z_e)}{z_e} = 0.\tag{2.38}$$

The root that falls within the desired range is

$$\begin{aligned}\alpha &= \frac{1}{3} \tan^{-1} \left[\sqrt{1 - \left(\frac{3}{2} \frac{z_u - (1 - z_e)}{z_e} \right)^2} / \left(\frac{3}{2} \frac{z_u - (1 - z_e)}{z_e} \right) \right] \\ \beta &= -\cos \alpha + \sqrt{3} \sin \alpha \\ z &= \beta z_e + 1 - z_e\end{aligned}\tag{2.39}$$

The axial velocity is transformed with the same method above except that the distribution function is used in place of the line charge. The distribution function that was derived in Neuffer's paper is

$$f(z, z') = C \sqrt{k^1(z_e^2 - (z - z_e)^2) - (z')^2}.\tag{2.40}$$

Here, z' is ratio of v_z and v_{beam} and C and k^1 are constants. The extrema of z' are given by Neuffer as

$$z'_{\text{max}} = -z'_{\text{min}} = \sqrt{k^1(z_e^2 - (z - z_e)^2)}\tag{2.41}$$

Figure 2.2 shows the distribution function versus z' for several values of z . The transformation for the end of the beam where $0 < z < z_e$ is done as follows. Setting the extrema of z' equal to ± 1 at $z = z_e$ gives $k^1 = 1/z_e^2$. Rewrite f as

$$f(z, z') = C\sqrt{1 - \beta^2 - (z')^2}. \quad (2.42)$$

with $\beta = (z - z_e)/z_e$. Integrating f over z' from z'_{\min} to z' gives

$$\int_{z'_{\min}}^{z'} f dz' = \frac{C}{2} \left[z' \sqrt{1 - \beta^2 - (z')^2} + (1 - \beta^2) \sin^{-1} \left(\frac{z'}{\sqrt{1 - \beta^2}} \right) + \frac{\pi}{2} (1 - \beta^2) \right] \quad (2.43)$$

Now, define f_u to be parabolic in z_u and uniform in z'_u over the range $0 < z_u < 1$ and $-\frac{1}{2} < z'_u < \frac{1}{2}$ such that the integral of f_u over all of z'_u is equal to the integral of f over all of z' .

$$f_u = C \frac{\pi}{2} (1 - \beta^2) \quad (2.44)$$

Integrating this over z'_u from $-\frac{1}{2}$ to z'_u gives

$$\int_{-\frac{1}{2}}^{z'_u} f_u dz'_u = C \frac{\pi}{2} (1 - \beta^2) (z'_u + \frac{1}{2}) \quad (2.45)$$

Now set equations 2.43 and 2.45 equal to each other and rearrange.

$$\frac{z' \sqrt{1 - \beta^2 - (z')^2} + (1 - \beta^2) \sin^{-1}(z'/\sqrt{1 - \beta^2})}{\pi(1 - \beta^2)} - z'_u = 0 \quad (2.46)$$

This equation is transcendental and has no analytic solution and so was solved iteratively by the Newton-Raphson method. The equation iterated is

$$z' = z' - \left[\frac{z' \delta + (1 - \beta^2) \sin^{-1}(z'/\sqrt{1 - \beta^2})}{(1 - \beta^2)\pi} - z'_u \right] / \left[\frac{2}{\pi} \frac{\delta}{1 - \beta^2} \right] \quad (2.47)$$

where $\delta = \sqrt{1 - \beta^2 - (z')^2}$. An initial guess is made based on z'_u .

$$z'_0 = \frac{\pi}{2} z'_u \sqrt{1 - \beta^2} \quad (2.48)$$

The iteration is done four times, giving acceptable convergence. The transformation in the end of the beam where $1 - z_e < z < 1$ is the same except the definition of β




Figure 2.2: Cigar velocity distribution function at several values of z . Velocity is normalized so the extrema are ± 1 and f is normalized to have a maximum of 1. $z = .5$ is in the center of the beam and the other two are nearing the end of the beam, which is at $z = 1$.

changes. $\beta = (z - 1 + z_e)/z_e$. Also, in the center of the beam, the transformation is the same, but $\beta = 0$.

The particle loader uses the square root of the normalized line charge to scale the transverse positions and velocities. The scaling is done in a way that keeps the tune depression in the ends of the beam the same as in the bulk (the ratio of the beam current and the transverse emittance, I/ϵ_x and I/ϵ_y , remains constant). The scaling is done in this way so that the envelope need only be calculated once. The beam envelope is obtained by scaling the calculated envelope by the square root of the line charge. This also keeps the space charge density, ρ , constant since

$$\rho = \frac{\lambda(z)}{Area(z)} = \frac{\lambda(z)}{a(z)b(z)\pi} \quad (2.49)$$

where $Area(z)$ is the transverse area of the envelope and a and b are the envelopes of the scaled beam in the x and y directions respectively. The transverse emittance varies in the same way as the line charge. The transverse velocity depends on both the scaled envelope and transverse emittance.

A finite length beam must be axially confined in some manner to prevent axial expansion from space charge and thermal forces. In a real ion induction accelerator, this is done by applying time-dependent axial fields periodically to the head and tail of the beam. They typically accompany the accelerating fields and are called “ears”. In WARP, this is generally modelled by applying the axial field continually in the beam frame. Two methods for calculating the applied field have been implemented. In both cases, the applied fields are functions of only z . In the first, the applied field is calculated directly from the beam space charge and emittance. The applied field is given by

$$E_{\text{ears}}(z) = \begin{cases} E_l(z - z_{\text{beam}} + z_s), & \text{for } z < z_{\text{beam}} - z_s; \\ 0 & \text{for } z_{\text{beam}} - z_s < z < z_{\text{beam}} + z_s; \\ E_l(z - z_{\text{beam}} - z_s), & \text{for } z > z_{\text{beam}} + z_s; \end{cases} \quad (2.50)$$

where z_{beam} is the lab frame location of the beam center, z_s is half the length of the center section of the beam, and

$$E_l = -\frac{2Ig}{2V_z z_e^2 \pi \epsilon_o} - \frac{MV_z^2 \epsilon_{\text{long}}^2}{Q z_e^4}. \quad (2.51)$$

Here, I is the current in the center of the beam, g is a geometric factor of order unity that depends on the size of the beam and beam pipe, V_z is the beam velocity, M is the mass of the simulation particles, ϵ_{long} is the longitudinal emittance, $\epsilon_{\text{long}} = 2z_e v_{\text{thz}}/v_b$, and Q is the charge carried by each of the simulation macro-particles. The longitudinal emittance is of the ends of the beam, $\epsilon_{\text{long}} = 2z_e v_{\text{thz}}/v_b$ where v_{thz} is the axial thermal velocity, and v_b is the beam velocity.

The “one-dimensional” assumption that the axial electric field is proportional to $\frac{d\lambda}{dz}$ is used to obtain this expression. This does not work well in three dimensions. The axial fields are not linear and the particles off axis are pushed too hard and bunch up where the taper meets the center portion of the beam.

The second method is to use the negative of the axial electric self-field on axis as the confining field. The field in the tapered ends of the beam is used as well as the field several grid points into the center portion for continuity. This method works

very well. Several runs have been done where the beam travels about one hundred times its length with very little distortion in the ends and no noticeable bunching. See chapter 3.

An example of a cigar shaped beam is shown in figure 2.3. The line charge density and various profiles are shown. Figure 2.4 shows a comparison of the electric field on axis to the confining field as calculated by both methods. The confining field calculated by the first method is linear and does not fit well. The slight difference in the axial electric field and the confining field calculated by the second method is due to the need to account for the thermal forces.

2.1.4 Calculation of Particle Moments

There are two different systems for calculating particle moments. One calculates moments versus z in the beam frame, the other calculates moments versus time in the lab frame. The beam frame moments are calculated directly from the particles on a one dimensional mesh. The lab frame moments are then calculated from the beam frame moments by interpolation.

For the beam frame moments, the particle data is loaded onto the moments mesh either by nearest grid point or by linear interpolation to the two nearest grid points. For linear weighting, errors caused by variations of particle axial positions within the finite-length grid cell are removed by “extrapolating” the particle quantities in time to the two grid points. Combining linear interpolation and this extrapolation decreases noise in the moments, since particles which are farther from the grid point and therefore have a worse extrapolation contribute less.

The extrapolation is done using the current and previous velocities of the particle. The change in time, Δt , from the current time and the time the particle was at the previous grid cell is calculated using the current axial velocity.

$$\Delta t = \frac{z_i - z}{v_z} \quad (2.52)$$

where z_i is the location of the previous grid point, z is the current location of the particle, and v_z is the current axial velocity. This time is used along with the current

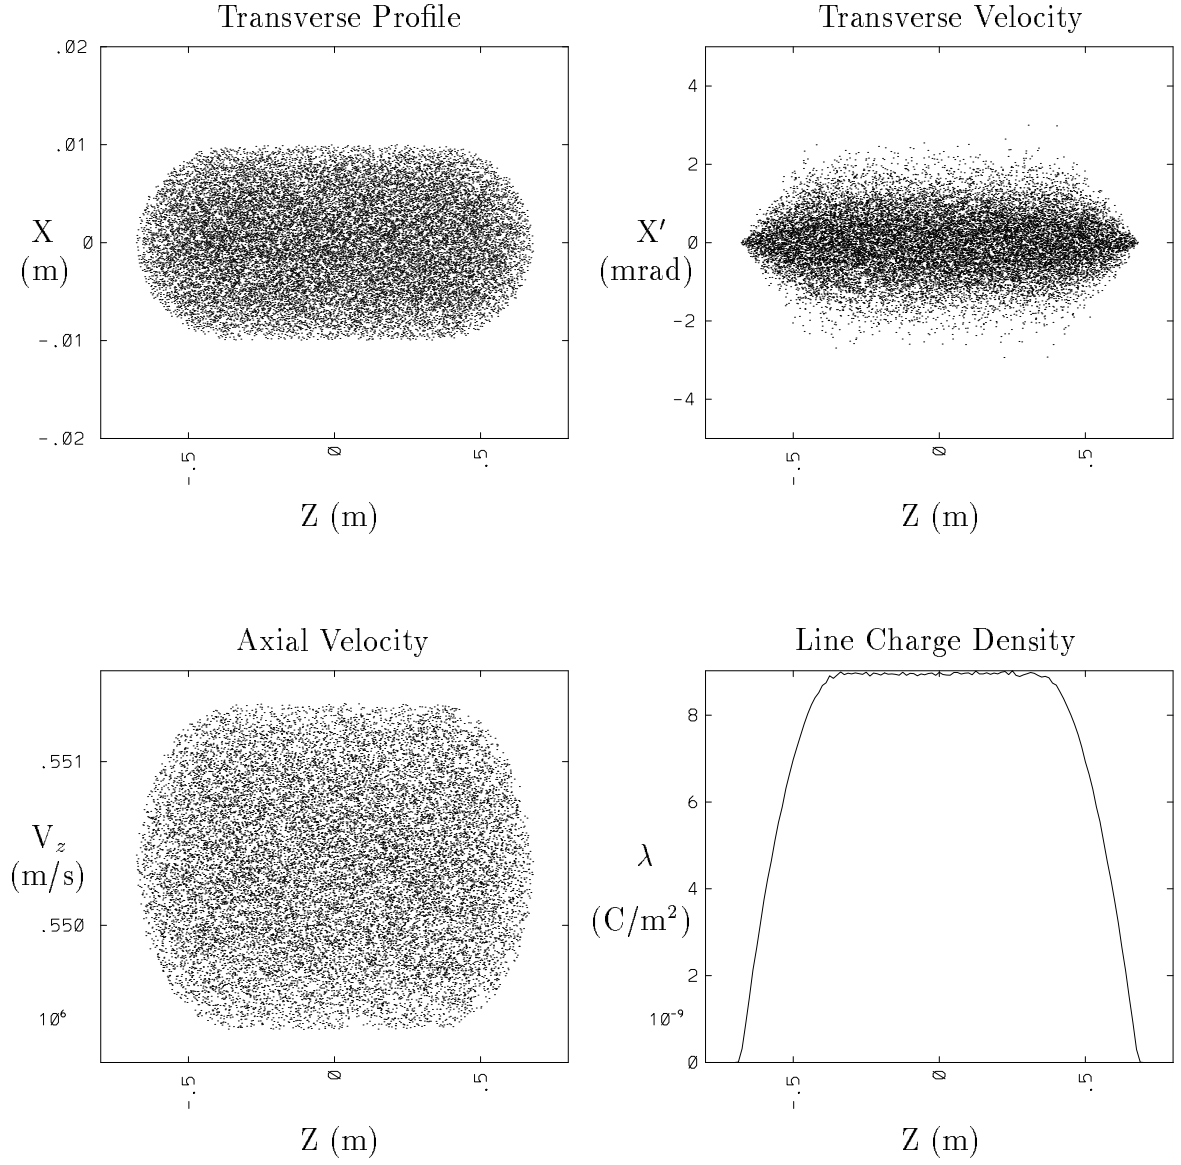


Figure 2.3: Cigar beam profiles and line charge density. Note that the falloff in λ is parabolic.




Figure 2.4: Comparison of axial electric and the negative of the ear fields on axis. The gap between the axial field and the ear field results from the addition of the field to counteract the thermal pressure.

transverse velocity, v_x (or v_y), and position, x (or y), to find the transverse position at the grid point, x_i (or y_i).

$$x_i = x + v_x \Delta t \quad (2.53)$$

The velocity is then linearly extrapolated using the previous and current velocity.

$$v_i = v \left(1 + \frac{\Delta t}{dt} \right) - v_{\text{prev}} \frac{\Delta t}{dt} \quad (2.54)$$

where v_i is the velocity at the grid point, v is the current velocity, v_{prev} is the previous velocity and dt is the time step between the current and previous velocity. The same equation is used for each of the velocities v_x , v_y , and v_z . A similar extrapolation is done for the other grid cell adjacent to the particle.

The main effect of the interpolation is to smooth the emittance calculation. In a quadrupole focusing element, the transverse velocities of the particles change rapidly. Particles with differing axial position but near a grid point inside of a quadrupole focusing element will have a wide variation in transverse velocity; this would cause a large increase in the calculated RMS emittance. Interpolation alleviates that increase by taking the contribution from the particles at the same values of z , when they have similar velocities. The interpolation assumes that any forces are constant throughout

the time step and so breaks down across the entrances and exits of quadrupole focusing elements. At an entrance or exit, residence corrections are used, so the focusing force is applied on only part of the time step. The extrapolation, assuming that the force was evenly applied throughout the time step, will then be incorrect. These errors are usually small, and can sometimes be seen in time-histories of the emittance as thin spikes.

2.2 Fields

The following sections describe the handling of the fields.

2.2.1 Beam Self-Fields

The fields from the beam space charge are calculated on a cubic three dimensional mesh. The fields are calculated either by fast Fourier transforms (FFT) or by successive over relaxation (SOR). The FFT's are sine-sine-periodic transforms. Using sine-sine transforms in the plane perpendicular to the direction of beam propagation forces the potential to vanish on the conducting walls. The axial transformation is periodic. The SOR field solver is described in detail in a later section.

2.2.2 External Fields

The focusing fields can be either alternating gradient (AG) quadrupole fields, or (as an idealization) continuous radial electric fields. The AG fields can be either magnetic or electric. The quadrupole fields in the simplest model are sharp edged with no fringing fields. Their profile is also flat; the fields have no axial dependence inside the element and are zero outside. The idealized continuous fields are constant along the axis.

The bending fields are from either magnetic or electric dipole elements. These are also sharp edged and have a flat axial profile with no fringing fields. (The lowest-order effects of slanted entry and exit faces are included for magnetic dipoles.) In a bend, each particle is tracked in its own coordinate system that rotates around the

bend.[22][19] The fields are also modified in a bend. Both the bends and dipoles can only affect the beam in the horizontal, or x , direction. There are magnetic dipoles B_y , with y being the vertical direction, and electric dipoles E_x . More detail on this can be found in section 2.2.3.

The accelerating fields as implemented have a flat profile within the gap and no fringing fields. They have no axial or transverse dependence. They are also constant in time (this restriction can be relieved by appropriate use of the code's interactive interface).

The focusing, bending, and accelerating fields are all set by the user with the lattice description formalism. This formalism is described in the next section.

2.2.3 Lattice Description

The lattice description consists of a set arrays for each element type that contain the starting and end points as well as other properties, such as field strength. The set of arrays gives all of the information needed about each element. Four element types are implemented: quadrupole focusing, dipole, bending (coordinate transformation), and accelerating gaps. The lattice can be periodic, either in the sense of repetitive elements along a linac, or in the sense of a circular machine.

The set of arrays for the quadrupoles hold the field strengths for either magnetic or electrostatic focusing and also the dimensions of electrostatic quadrupoles. These dimensions are used by the more elaborate capacity matrix and SOR field solvers to model the electrostatic quadrupoles from first principles.

The quadrupole elements can be made more realistic by adding offsets representing misalignments or by adding non-linear, higher order fields. The offsets can be random, with a Gaussian distribution, or can be set by the user. The periodicity can be the same as that of the general lattice, or can be different. The offset would have the same periodicity in a recirculating accelerator, when the beam would repeatedly pass through the same elements. For a straight linac, they could be different; the lattice would be periodic, but the beam would be passing through elements with different

errors. Some higher order moments are included along with focusing fields. They are the dodecapole component (the V_{66} term in equation 1.4) for the sharp-edged electrostatic focusing model and the octopole and octopole-like fields (the term V_{40} , V_{42} and V_{44} in equation 1.4).

The bend lattice is usually tied to the dipole lattice. If the user only sets the bend lattice, a magnetic dipole lattice will be set automatically to give the appropriate bends. The bending elements are then assumed to be dipole magnets with the strength determined by the bend radius and beam parameters:

$$B_y = \frac{M_{\text{ion}} V_{\text{beam}} m \bar{\gamma}}{Z_{\text{ion}} r_b} \quad (2.55)$$

where M_{ion} is the mass of the beam ions, V_{beam} is the beam velocity, $\bar{\gamma}$ is the average relativistic factor of the beam, Z_{ion} is the charge of the ions, and r_b is the bend radius. The starts and ends of the dipole elements are the same as the bend elements. Of course, the user can manually over-ride this automatic association.

2.2.4 Internal Lattice Description

Inside the code, the lattice described by the user is transformed into an internal representation that is more efficient. For each of the lattice arrays, there is an internal array that holds the same data, but in a different format. The internal arrays are one-dimensional grid arrays along the axis where each grid cell contains information about the nearest element. The starts and ends of the elements are converted to the lab frame. These arrays move with the beam and are set before every time step.

Obtaining the nearest element to each particle is thus efficient. Its axial position determines the grid cell where the particle lies. The nearest element is then immediately known from the internal lattice arrays. This means that scans through the input lattice description are done for each axial grid cell location, of which there are few, instead of for each particle, of which there are many.

2.2.5 Field Solutions with Internal Conductors

Two different types of internal conductors can be handled in the electrostatic field solution. A structure that is infinite in extent along the z axis and has no z dependence, such as a round pipe, can be modeled, as can periodic electrostatic focusing elements. The first is dealt with via the capacity matrix method [5] which is built on top of the 3-D Fast Fourier Transform field solver. The second can be dealt with either via the capacity matrix method, or, for more complex focusing structures, the point successive overrelaxation (SOR) method [11].

Capacity Matrix Field Solver

For the pipe-like structures, the capacity matrices are applied in k_z space, with each wave number treated independently. This can be done with the assumption that there is no variation in z of the conductor cross section. Each wave number then has its own matrix. The inverse matrices are found first by using $C^{-1}\rho_{k_z} = \phi_{k_z}$, where ρ_{k_z} is the charge density at the points on the conductor in k_z space, and ϕ_{k_z} is the potential at the points on the conductor in k_z space. Both ρ_{k_z} and ϕ_{k_z} are one dimensional arrays with an element for each point on the conductor in the transverse plane. For each such point, the charge density for all k_z is set to unity and the electrostatic potential is found with the 3d FFT without the forward and backward z transforms. The array ϕ_{k_z} fills the appropriate column (or row, as it is symmetric) of the inverse capacity matrix. Once the inverse matrix has been completely found, it is inverted to find the capacity matrix. The inversion is done by first reducing the matrix by Gaussian elimination and then calculating the inverse column-by-column via backsubstitution. The matrix is then used in k_z space by solving the equation $\rho_{k_z}^{\text{induced}} = C\phi_{k_z}$. The $\rho_{k_z}^{\text{induced}}$, the induced charge on the conductor, is loaded onto the grid, along with the z transformed beam space charge. The potential is then solved for again via the three dimensional FFT, including both the beam space charge and the induced charge.

All of this is done assuming an eightfold symmetry of the conductor shape within the square walls of the FFT. This reduces the amount of work to calculate the capacity

matrix and the amount of storage for the matrix by a factor of nearly eight, though it makes the coding quite complicated.

The periodic electrostatic focusing elements are handled completely in 3d real space. For quadrupole rods which are far apart in z , the capacity matrix method is used. Since the elements are far from each other, they can have independent matrices. There is otherwise no difference from the standard capacity matrix method. Each matrix represents the four symmetric conductors that make up the quadrupole. This allows use of fourfold symmetry to save space and computation time.

Quadrupoles are considered far apart when their separation is greater than the aperture. Field solutions were found with various configurations to examine the interaction between neighboring quadrupoles. The gauge of the error is the ratio of the calculated potential on the conductor and the desired potential. With increasing interaction between quadrupoles, the potential calculated will deviate further from the desired potential. The length of the separation was varied from .01 to .1 meters. The quadrupole aperture was varied from .002 to .2 meters. The transverse extent of the field mesh was varied transversely to cover just under half of the quadrupoles to containing the entire quadrupole. Finally, the length of the quadrupole was varied. Table 2.2.5 shows some of the data taken. The dominant effect is the ratio of the separation to the aperture. The numbers were used to derive the following relation.

$$r = 1 - 10^{-\gamma \frac{s}{a}} \quad (2.56)$$

Here, r is the ratio of potentials, s is the separation, and a is the aperture. The parameter γ is of order unity. For the typical cases, it varies between .8 and 1.2. For the more extreme cases, γ varied from .25 to 2.

Care is needed in choosing the location of the conductors with respect to the grid. To avoid jagged fields near the conductor surface, the points used in the matrix are independent of the grid and conform to the surface of the conductor. The tri-linear interpolation is used to lay down the induced charge of each point onto and to obtain the potential at each point from the surrounding points.

The relation between transverse surfaces and the grid is subtle since the surfaces

	a=.02 wall=.05		a=.02 wall=.04		a=.2 wall=.4		a=.002 wall=.004	
s	r	γ	r	γ	r	γ	r	γ
.10	.99992	.819	.999991	1.009	.6917	1.022	$1. - 3 \times 10^{-13}$.25
.08	.99946	.816	.999904	1.004	.6152	1.037		
.06	.9964	.814	.999006	1.001				
.04	.976	.809	.9897	.994	.4057	1.130		
.02	.843	.804	.8964	.985				
.01	.607	.811	.6771	.982	.1671	1.588	.9928	.43

Table 2.2: Check of interaction of quadrupoles with capacity matrix field solver. r is the ratio of the potential on the conductor obtained in the field solve and the desired potential. The data was fit to $1 - 10^{-\gamma \frac{s}{a}}$. a is the quadrupole aperture (the rod radius is $8/7a$). The wall is the extent of the field mesh.

are curved and since accurate focusing fields are needed. The points do not lie on the grid so if there are fewer points than grid points, there will be holes through which the field can leak. The focusing fields are slightly higher than what would be expected. With fewer points on the conductor surface, more charge is needed at each point to give the desired voltage. This increased charge is deposited on grid cells surrounding the conductor point, putting some of it closer to the axis, decreasing the effective aperture. The focusing field is thus increased. The difference is typically on the order of a few percent or less. Figure 2.5 shows values of the ratio of the peak focusing field to the expected focusing field and the integrated focusing field versus number of points per axial plane in the conductor. These numbers will vary depending on the grid cell size in relation to the conductor size. Note that after a certain number of points, the accuracy does not improve any further.

Along the axis, the points are separated by the grid cell size. The precision of the field solution is very insensitive to the location of the point with respect to the grid. Field solutions with the points in different places with respect to the grid were compared. The integral of the resulting quadrupole focusing field over one quadrupole was the same out to ten places for all of the cases!

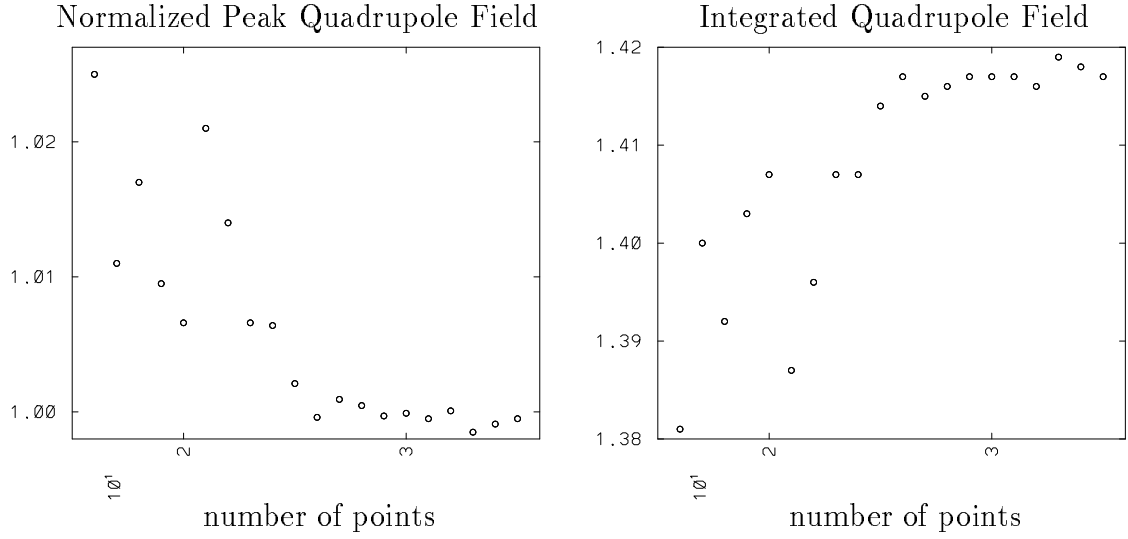


Figure 2.5: Focusing fields versus number of conductor points. a) is the ratio of the peak quadrupole fields and the expected quadrupole field. b) is the integral of the quadrupole field over one quad in arbitrary units. In both cases there is a wide spread of values for small numbers of points. When the number of points is odd, there is a point at the tip of the rod. When the number is even, there is no point at the tip and the rod is effectively farther from the axis.

Successive Overrelaxation Field Solver

For more complex structures that are close and possibly touching, like long quadrupoles with end plates, the capacity matrix may become too big for solution in a reasonable amount of time. The matrix becomes bigger when the structures take up more space inside the field grid and require more points to be described. Also, the elements cannot be treated individually so all must be included in one very large matrix. For cases like these, the point SOR method was competitive and therefore implemented. Currently, the code is set up to specifically handle quadrupole focusing structures with end plates. See figure 6.1 for a diagram showing the geometry.

In three dimensions, the finite difference form of Poisson's equation used is the seven point difference scheme. After rearrangement and inclusion of the relaxation parameter, the equation that is iterated for each grid point is:

$$\begin{aligned} \phi_{ijk} = & \omega \Delta^2 (\rho_{ijk} / \epsilon_0 + \frac{\phi_{i+1jk} + \phi_{i-1jk}}{\Delta x^2} + \frac{\phi_{ij+1k} + \phi_{ij-1k}}{\Delta y^2} + \frac{\phi_{ijk+1} + \phi_{ijk-1}}{\Delta z^2}) \\ & + (1 - \omega) \phi_{ijk}. \end{aligned} \quad (2.57)$$

where the subscripts are grid cell indices, ω is the relaxation parameter,

$$\frac{1}{\Delta^2} = \frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2}, \quad (2.58)$$

and Δx , Δy , and Δz are the grid cell sizes in the various directions. The ordering that is used is the three dimensional extension of even-odd. The potential on the internal conductors is enforced by setting it to the desired value before each iteration. The boundary conditions at the edge of the grid can be either Neumann, Dirichlet, or periodic.

For simple SOR, all of the conductor points are on the grid. This gives jagged fields near the conductors and also gives incorrect focusing fields. When only points within the conductor are considered to be at the fixed potential of the conductor, the focusing field on axis will be reduced by as much as five to ten percent. The effective aperture is approximately given by the distance from the centerline to the grid points inside the conductor which are nearest the surface. These points will always be farther away from the centerline than the actual conductor surface and so the fields on axis will be weaker.

There are two solutions to that problem. The first is to change the points which make up the conductor by including points just outside the surface. The optimum distance outside the surface over which to include points was empirically determined with several different grid cell sizes and several conductor shapes to be one quarter of the grid cell size. The better way is to explicitly include sub-gridscale boundary placement into Poisson's equation. This was done and is described in section 2.2.5.

The point SOR proved slow for simple internal structures. The main problem with the iterative method is that all of the work must be redone at at each time step. The capacity matrix is superior for simple internal structures since it does most of its work at the start of the run when it calculates the matrix. This leaves only a small amount of work at each time step, the matrix multiply and extra 3D FFT. The iterative solution can be sped up by using the solution from the previous time step, but the fields change enough between time steps so that this does not help enough to render SOR faster than the capacity matrix method when the quadrupoles can be

treated independently.

Another iterative method that would include internal conductors was also examined. It is an iterative method based on the algorithm developed in reference [12]. This method was found to be far too slow for our purposes. Each iteration required a full 3d FFT field solution. The method would be appropriate were resistive blocks desired.

Sub-gridscale Boundary Conditions

Sub-gridscale placement of boundaries is handled by explicitly including the location of the boundary in relation to the grid in the finite difference form of Poisson's equation. This is easily explained in 1-D. The finite difference form of Poisson's equation, $\nabla^2\phi = -\rho/\epsilon_0$, is

$$\frac{\phi_{i+1} - 2\phi_i + \phi_{i-1}}{\Delta^2} = -\frac{\rho_i}{\epsilon_0}, \quad (2.59)$$

where the subscript refers to the grid cell index and Δ is the grid cell size. If the edge of the conductor surface is between the grid points i and $i+1$, and $i+1$ is inside, the potential at the edge would be interpolated as

$$\phi_{\text{edge}} = (1 - \delta)\phi_i + \delta\phi_{i+1}. \quad (2.60)$$

Here, δ is the distance between the grid point i and the surface divided by the grid cell size. The value of ϕ_{edge} is the known voltage on the conductor, and the value ϕ_{i+1} is a free parameter since it is inside the conductor. The previous equation can be rearranged to give a value to ϕ_{i+1} .

$$\phi_{i+1} = \frac{\phi_{\text{edge}} - (1 - \delta)\phi_i}{\delta} \quad (2.61)$$

This is then put back into Poisson's equation to give

$$\frac{(\phi_{\text{edge}} - (1 - \delta)\phi_i)/\delta - 2\phi_i + \phi_{i-1}}{\Delta^2} = -\frac{\rho_i}{\epsilon_0}, \quad (2.62)$$

which is then rearranged to bring all of the terms of ϕ_i to the left hand side to make it fully explicit.

$$\phi_i = \frac{\phi_{i-1} + \phi_{\text{edge}}/\delta + \Delta^2\rho_i/\epsilon_0}{2 + \frac{1-\delta}{\delta}} \quad (2.63)$$

Note that in the limit $\delta \rightarrow 0$, ϕ_i approaches ϕ_{edge} . The expression is never used though in the code when δ is zero.

The SOR relaxation parameter, ω , is added in afterward to give the final equation to be iterated:

$$\phi_i = \left(\frac{\phi_{i-1} + \phi_{\text{edge}}/\delta + \Delta^2 \rho_i/\epsilon_0}{2 + \frac{1-\delta}{\delta}} \right) \omega + \phi_i(1 - \omega). \quad (2.64)$$

The parameter ω has the same meaning as in the previous section.

The idea extends easily to two and three dimensions. The linear interpolation is done independently in each direction for which there is a conductor. From each of the interpolation equations, the potential at the point inside the conductor is put in terms of the point on the surface and the point outside. The expressions are then each put into the seven point finite-difference form of Poisson's equation which is then rearranged to make it fully explicit. If the one point ϕ_{i+1jk} is in the conductor, equation 2.61 (adding the indices j and k to ϕ) is used to replace ϕ_{i+1jk} in the finite-difference equation about the point ϕ_{ijk} .

$$\begin{aligned} & \frac{(\phi_{\text{edge}} - (1-\delta)\phi_{ijk})/\delta - 2\phi_{ijk} + \phi_{i-1jk}}{\Delta x^2} + \frac{\phi_{ij+1k} - 2\phi_{ijk} + \phi_{ij-1k}}{\Delta y^2} + \\ & \frac{\phi_{ijk+1} - 2\phi_{ijk} + \phi_{ijk-1}}{\Delta z^2} = -\rho_{ijk}/\epsilon_0 \end{aligned} \quad (2.65)$$

The resulting equation is rearranged to make it explicit and the relaxation parameter, ω , is added in.

$$\begin{aligned} \phi_{ijk} = & \frac{\omega \Delta^2 (\rho_{ijk}/\epsilon_0 + \frac{\phi_{\text{edge}}/\delta + \phi_{i-1jk}}{\Delta x^2} + \frac{\phi_{ij+1k} + \phi_{ij-1k}}{\Delta y^2} + \frac{\phi_{ijk+1} + \phi_{ijk-1}}{\Delta z^2})}{1 + \frac{(1-\delta)\Delta^2}{\delta \Delta x^2}} \\ & + (1 - \omega)\phi_{ijk}. \end{aligned} \quad (2.66)$$

Now, during the iterations, this equation is used at the point ijk instead of the equation 2.57.

When there are two points inside the conductor in different directions from point i, j, k , they can be treated independently. The interpolations are done for each of the two points and then put back into the finite difference equation. If the points ϕ_{i+1jk}

and ϕ_{ij+1k} are inside the conductor, the final equation would be

$$\phi_{ijk} = \frac{\omega \Delta^2 (\rho_{ijk}/\epsilon_0 + \frac{\phi_{\text{edge}}/\delta_x + \phi_{i-1jk}}{\Delta x^2} + \frac{\phi_{\text{edge}}/\delta_y + \phi_{ij-1k}}{\Delta y^2} + \frac{\phi_{ijk+1} + \phi_{ijk-1}}{\Delta z^2})}{1 + \frac{(1-\delta_x)\Delta^2}{\delta_x \Delta x^2} + \frac{(1-\delta_y)\Delta^2}{\delta_y \Delta y^2}} + (1 - \omega)\phi_{ijk}. \quad (2.67)$$

Here, δ_x and δ_y are the respective unitless distances in the x and y directions between the grid point ijk and the surface. The extension to three dimensions is done similarly.

Note that currently, the points must be offset from ijk along differing axes. The case when two opposite points, say ϕ_{i-1jk} and ϕ_{i+1jk} , are inside the conductor, while ϕ_{ijk} is outside, has not been implemented. This limits the use of the sub-gridscale coding to convex or planar surfaces. For example, with a plate in which there is a hole, the sub-gridscale coding can not be used. Only the potential on grid points within the plate can be set, leaving the inside of the hole jagged. Having an end plate with a jagged edge has little effect on the simulation. The plates are narrow (in z) so the jagged edge is over only a small extent.

Two simulations were done to compare the results with the sub-gridscale boundary conditions and without. The ESQ injector experiment which is discussed in Section 6.5 was simulated. The results are shown in figure 2.6. The envelope without the sub-gridscale boundary conditions was larger since the quads are effectively farther away from the axis. The emittance is also subsequently larger since the larger envelope experiences larger non-linear fields. Another simulation was done in which the aperture of the plates holes was varied. There was no noticeable change in the results, showing that the results are insensitive to the size of the plate aperture and so it is not necessary to resolve the plate aperture with the sub-gridscale boundary conditions.

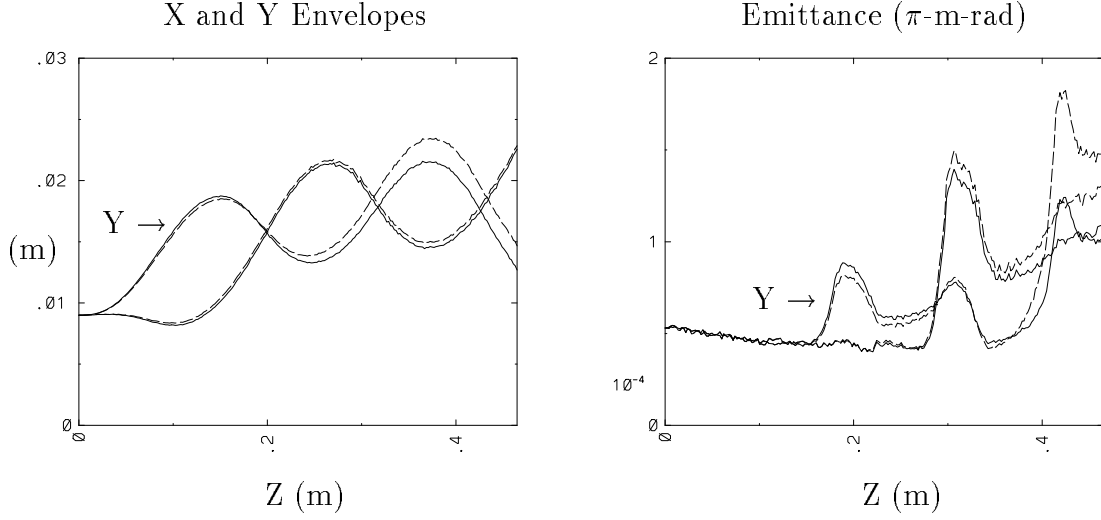


Figure 2.6: Comparison of the envelope and emittance of an ESQ injector with and without sub-gridscale boundary conditions. The solid line is with and the dashed line is without. Note that the dashed envelope is larger since without sub-gridscale boundary conditions the quadrupoles are effectively farther away from the centerline and so the focusing force is less. The dashed emittance is also larger since the larger envelope experiences larger non-linear fields.

2.3 Optimizations, Plot Handling, and the Basis System

2.3.1 Techniques Used for Optimization

Various techniques have been used to decrease both computational time and the amount of memory used. Some have been explained earlier, in the appropriate sections, others are listed here.

In order to reduce the number of large three dimensional grid arrays, we do not store the electric fields on the grid. Instead, the potential in the neighborhood of each particle is gathered on a particle-by-particle basis. This ends up saving us both memory and time. Storing the fields on a grid would require an additional three grid arrays, one for each coordinate direction. Since the beam usually occupies the center region of the mesh, all of the outer areas would be wasted space. Also, the finite differencing of the potential to get the field in the outer areas would be wasted time.

Usually it is the case that there are more grid points than there are particles, so fewer total calculations are done on a particle by particle basis.

The electric fields are calculated as a two point finite difference, $E_i = \frac{\phi_{i-1} - \phi_{i+1}}{2\Delta x}$, where E and ϕ are the electric field and the electrostatic potential and i is the grid index in the appropriate direction. Each of the three fields, E_x , E_y , and E_z , are calculated on the eight vertices of the grid cell that the particle occupies and these are interpolated linearly to the particle. All field components are calculated and interpolated at the same time for each particle. This requires gathering the potential from thirty-two cells in the neighborhood of each particle.

A significant speedup can be gained by using a special feature of some computers that is known as “vectorization”. This allows simple loops to be processed at a much faster rate by overlapping the calculations. Loops that do any accumulation, though, cannot be vectorized. If accumulations are overlapped, some of the passes through the loop will not be added in, leading to errors. There are two major pieces of the code that do accumulation, the gather of the charge density from the particles, and the calculation of the moments from the particles which was described earlier. In a loop over the particles, more than one particle will be contributing to the same field or moment array element, leading to possible collisions when vectorization is employed. This is circumvented by vectorizing with respect to other quantities. For the charge density scatter summation, each particle will contribute to the eight grid points on the corners of the grid cell it occupies. These corners can be vectorized over. The moments calculations are accumulations of the positions and velocities of the particles. The list of moments that are calculated is the quantity that is vectorized over.

2.3.2 Plot Handling in WARP

The main form of output in WARP are plots created during a run and plots made at the end of a run. There are several different type of plots that can be made automatically during a run: particle plots, contour plots, and line plots. The particle plots

are plots of the simulation particles in various views. The particles that are included in the plot are determined either by subsets or by windows. A subset consists of a smaller number of particles than the total number, chosen randomly. The windows are ranges in space; all particles within a window are plotted. The contour plots are of the quantities on the grid, the charge density and the electrostatic potential. These are also done by windows (the plane at mid-window is plotted). The line plots are plots of various quantities versus z . The plots made at the end are the history plots. Plots of any type and of any quantity can be made through the Basis interpreter, which is described briefly in the next section.

2.3.3 Basis

Basis[13] is a code development and run-time system that, through an interactive interpreter, gives the user access to variables and subroutines in the compiled code. The interpreter also has Fortran-like constructs (including loops and conditional statements) to manipulate variables. These, combined with the access to code variables, allows the user to dynamically control the progress of a run, modify internal variables, and produce diagnostic plots and output in an interactive manner. A file containing assignments to variables and calls to initialization routines can be read in to set up and start a run; this is a more sophisticated version of the Fortran “namelist” capability, but allows the user to “reprogram” the code’s behavior by including the necessary coding in the input deck. Basis also takes care of memory management and allows dynamic allocation of arrays, alleviating the need to recompile in order to change array sizes.

Part II

Applications

Chapter 3

Long Time Behavior of Finite Length Beams

The study of long time beam behavior requires the capability of loading a finite length beam and confining it axially. Unconfined finite beams exhibit free expansion longitudinally due to both thermal forces and space charge forces. Beams run with an axial (z) velocity “tilt” (head to tail gradient) are initially compressed, but expansion soon takes over again, blowing the tips off the beam. The capability of loading and confining cigar-shaped, finite length beams has been implemented in WARP and is described in section 2.1.3.

Many runs have been carried out using the axial confinement model; we summarize some of them here. See table 3 for the parameters of the runs described. Run number one, initialized with the parabolic ends each one quarter the total length of the beam, uniform focusing, and cold in z , was run for 120 m. Another, run number two, was started off with the thermal v_z the same as the thermal v_\perp , the other parameters being the same, and was run for 60 m. Run number three was taken the farthest, to 210 m. It used magnetic quadrupole focusing with a tune depression of $\sigma_0 = 60^\circ$ to $\sigma = 20^\circ$ (phase advance per lattice period)[14], and the thermal v_z almost a factor of 10 less than the thermal v_\perp . All of these runs were well-behaved. The axial spreading was much less than 1% of the beam length, and less than 10 out of about 50000 particles

run number focusing ^a	1 uniform	2 uniform	3 quadrupole	4 quadrupole
tune ^b	60° to 20°	60° to 20°	60° to 20°	60° to 55°
a0 ^c (cm)	2.34	2.34	3.05	1.87
b0 ^d (cm)	2.34	2.34	1.88	1.12
current (Amps)	5.22	5.22	5.22	0.343
wall location (cm)	±5	±5	±9	±9
beam length (m)	1.2	1.2	1.2	1.2
taper len over total len	1/2	1/2	1/2	1/5
initial RMS vz (m/s)	50.	2×10 ⁴	4×10 ³	1×10 ⁵
run length (m)	120	60	210	144

^aUniform focusing means a radial electric field continuously applied to a cylindrical beam. Quadrupole focusing means strong focusing with magnetic quadrupoles. The quadrupole focusing lattice is periodic with periodicity of 1.2m. It is a FODO arrangement with magnets .2m long and a .4m drift space.

^bphase advance per lattice period, (undepressed) to (depressed); for uniform focused runs, the phase advance is measured as if there were a 1.2m period.

^cThe maximum initial semi-axis of the beam

^dThe minimum initial semi-axis of the beam

Table 3.1: Parameter list of runs used to study long time behavior of beams. All of the runs had the following parameters: $\epsilon_{\perp} = 16\text{mm-mrad}$, $V_{\text{beam}} = 1.27 \times 10^7 \text{m/s}$, ≈ 50000 simulation particles and a computational mesh with dimensions $64 \times 64 \times 128$.

were lost out the ends (because the confining fields do not extend beyond the initial tips of the beam). Also, the transverse emittance remained roughly constant in time over much of the beam, with no greater than 10% fluctuations; there was no net transverse heating in the center section, and only heating by a factor of 1.5 to 2 in the ends, signifying that the tapered ends were not initially in equilibrium. The transverse profile as well as an overlay plot of λ versus z at various times are shown for run number two in figure 3.1

Run number three is shown in more detail in figures 3.2 and 3.3. Plots of the x and y profile at the start and at the end of the run are shown in figure 3.2. The only differences are the edge of the beam which is somewhat ragged after the run, and the few particles that have escaped the head and tail of the beam. In figure 3.3, the transverse emittance, $\epsilon_{xx'}$, versus time at four places in the beam is shown. The only

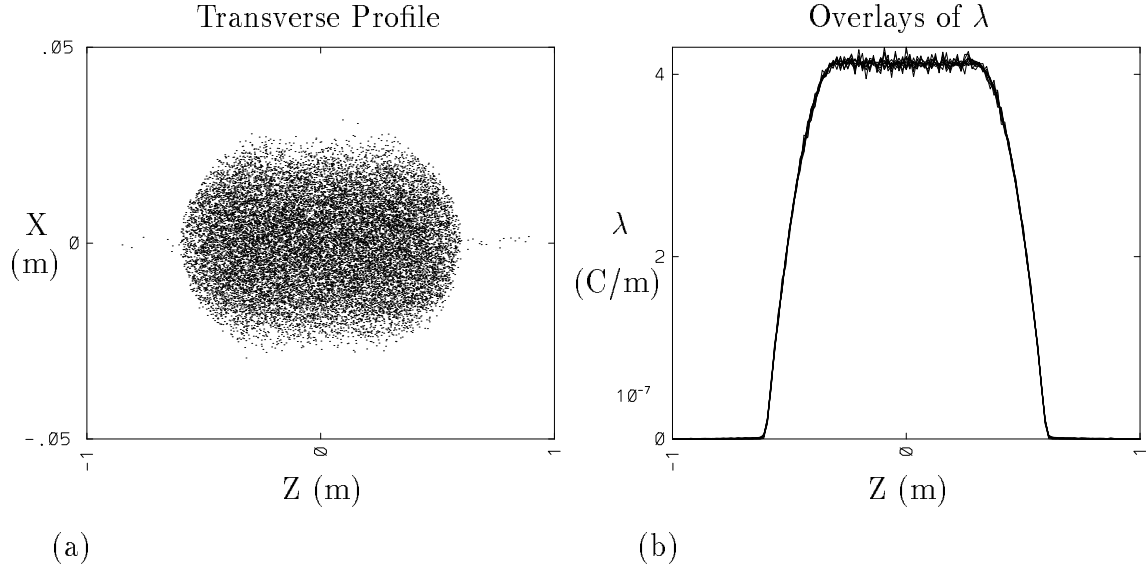


Figure 3.1: Uniformly focused cigar beam after 120 m. a) Transverse profile b) overlay of λ at various times.

one that shows growth is the one near the tip of the taper. The distribution used to initially load the ends of the beam is not an equilibrium distribution so the ends redistribute themselves as they approach an equilibrium with the result of emittance growth. Figure 3.3 also shows the z thermal energy. The initial rapid rise is an equilibration effect common to runs initialized colder in z than in x and y , which we believe to be physical[15]. The subsequent slower rise is the result of numerical heating common to particle codes; it is reduced by increasing the number of particles as shown on figure 3.3d. The z thermal energy and the transverse emittance in the tapers are the only parameters that show significant change.

A run that was emittance-dominated in both the transverse and longitudinal direction was also made (number four in table 3). The tune depression was 60° to 55° , and the thermal v_z was 50% larger than the thermal v_\perp . The model for the confining field was still accurate enough for this case; no waves were launched into the center of the beam. The beam was well confined. See figures 3.4 and 3.5.

The confining field is intended to model the confining fields used in a real accelerator. In the latter, the confining field is applied only at the acceleration gaps,

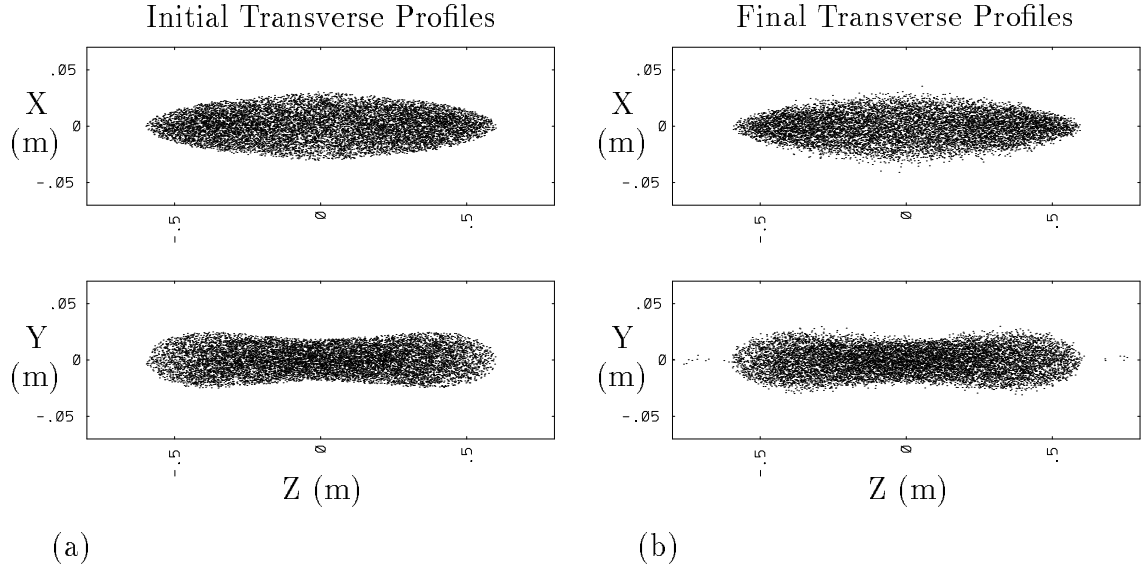


Figure 3.2: Initial and final (after 210 m) transverse profiles of an axially confined beam. The major visible differences are a more ragged beam edge and the few particles that have escaped the confining fields.

along with the accelerating fields; this is done by applying “ears” to the accelerating pulses. The true intermittently applied field is more difficult to derive since it must also compress the beam slightly to account for the free expansion between accelerating gaps. This more complicated field will have space and time dependence. Our initial efforts at modelling intermittent application of the confining force did not use a carefully tuned field, and led to large fluctuations. One note is that after several applications of the confining field, the beam did settle down into a new distribution, with moderate emittance growth. The work is being pursued by others[15]

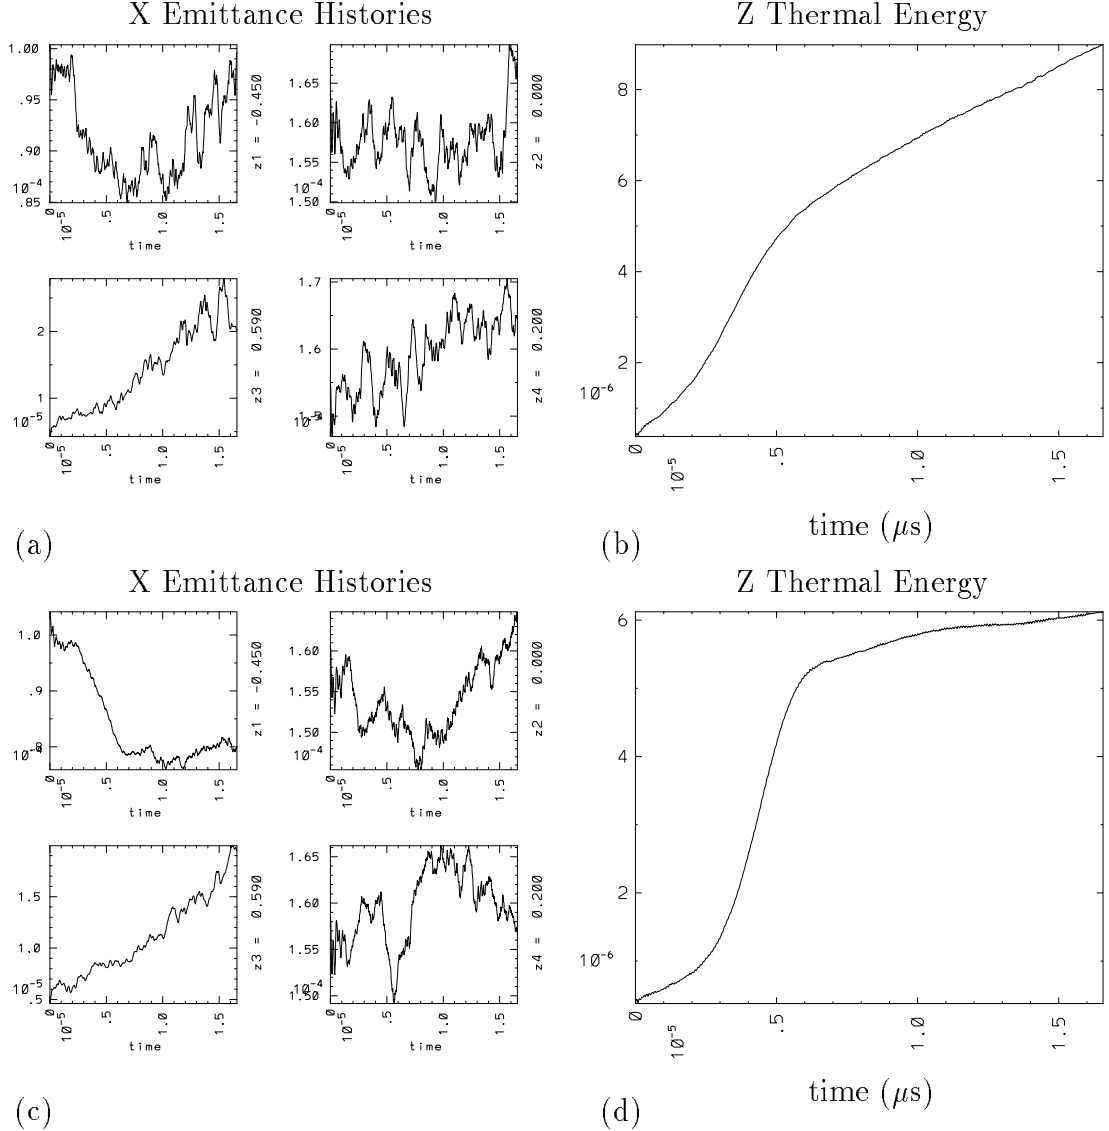


Figure 3.3: Emittance time histories of an axially confined beam. a) shows the emittance histories at various places along the beam. Most are fairly steady, except the emittance near the end of the beam, $z=.59$ m, which shows steady growth. The distribution used to initially load the ends of the beam is not an equilibrium distribution so the ends redistribute themselves as they approach an equilibrium with the result of emittance growth. b) shows the axial thermal energy history. The initial rapid rise is the result of a collective effect causing the equilibration of the transverse and axial temperatures. The latter slower rise is from numerical heating and can be avoided by increasing the particle number. c) and d) show the same quantities but with the number of particles increased by four. Note that the rise in the axial thermal energy in the second half of the run has become much more shallow.

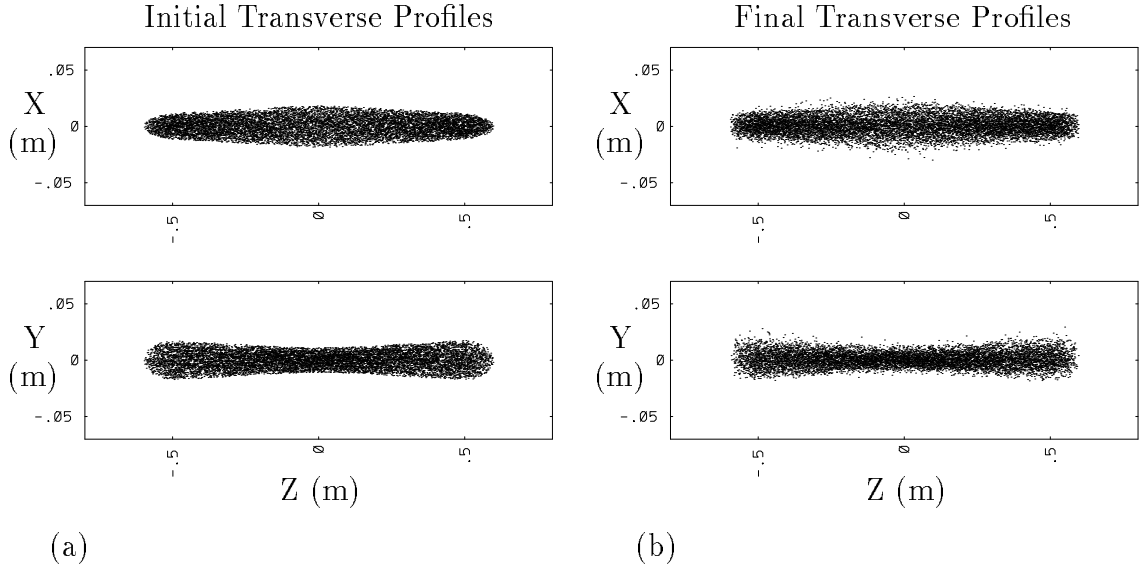


Figure 3.4: Initial and final (after 60 m) transverse profiles of an emittance dominated, axially confined beam. The major visible differences are a more ragged beam edge and the shape of the ends. The ends evolved into a blunt shape.




Figure 3.5: Emittance histories for an emittance dominated beam. The only significant change is near the ends of the beam, the plot in the lower right.

Chapter 4

Effects of Round and Square Pipes

The shape of the pipe surrounding the beam will affect the behavior of the beam. The image charges in different shape pipes will not be the same since the parts of the pipe that are closest to the beam and therefore have the largest image charge will not be the same. The shape of the surrounding pipe will have least effect on a beam which is centered on axis. The walls are relatively far away from the beam so the image charges on the walls will have the least effect. Off-axis beams will feel the effects of the pipe shape more strongly.

This section compares the behavior of beams in round and square pipes. This comparison is of interest in applications of this code since it works in Cartesian coordinates. The simplest boundaries are then the sides of a rectangular metal box. With the transverse dimensions the same, this gives a square beam pipe. In general, though, the true beam pipe shape would be cylindrical. The fields in the square pipe are calculated using fast Fourier transforms in Cartesian coordinates with transverse conducting boundaries. The fields in a round pipe are calculated using the capacity matrix method on top of the FFT in Cartesian coordinates as described in section 2.2.5.

Several axially confined beams were run through randomly offset quadrupole magnets. The offsets represent alignment errors. The offsets were all in the transverse plane, with x and y RMS values of 1 mm. The offsets move the beam off center;

a beam with maximum radius of 3 cm was shifted at most 3cm off axis in x and y when followed for 54 m. Runs with the walls at ± 5 cm and ± 7 cm were done with both round and square pipes. The runs had a tune depression of 60° to 20° , with the thermal v_z the same as the thermal v_\perp . Other parameters were: a beam major radius of 3.05 cm, a beam minor radius of 1.88 cm, a current of 5.22 Amperes, a beam length of 1.2 m, a ratio of taper length to total beam length of $1/5$, an initial RMS V_z of 2×10^4 m/s, an ϵ_\perp of 16 mm-mrad, and a V_{beam} of 1.27×10^7 m/s. The beam travelled through a lattice of magnetic focusing quadrupoles with a half lattice period of .6 m, quadrupole length of .2 m, and a focusing strength of 14.9 Tesla/m. The number of simulation particles was 50000 and the grid cell size was $64 \times 64 \times 128$. The runs took about 22 minutes on the NERSC C-90.

The fluctuations of the beam centroid depended only very weakly on the pipe shape. Figure 4.1a compares the centroid motion of beams in the two pipes. The biggest difference is near the end when the beam is farthest off axis. The beam in the round pipe is pulled slightly farther off axis, since, on average, the pipe is closer to the beam than the square pipe. There was also little difference in the fluctuations of the emittance as shown in figure 4.1b. These runs give us confidence that the use of a square pipe for efficiency will generally not have a deleterious affect on future applications. When the highest accuracy is needed, the more expensive round pipe field solver can be used.




Figure 4.1: Comparison of beams in round and square pipes. a) Centroid position, the beam in the round pipe is pulled slightly farther off axis because the pipe is closer to the beam, b) X emittance.

Chapter 5

Drift Compression on MBE-4

The MBE-4 experiment is a heavy ion accelerator that was designed to examine various aspects of the propagation and acceleration of space charged dominated heavy ion beams for heavy ion fusion. One series of experiments studied the effects of axially compressing the beam by imposing an initial head to tail gradient in the axial velocity. The results showed anomalous increases in the transverse emittance, increasing with the compression factor. Theory suggested that a compressing beam in pure quadrupole fields should not experience any emittance growth. The WARP3d code was used to simulate these drift compression runs. The effects of various initial distributions, non-linear focusing fields, and image charges have been examined. The dodecapole field and other non-linear fields had the largest effect toward increasing the emittance during compression. (In this chapter, the emittance referred to is always the unnormalized emittance.)

5.1 MBE-4 Results

After the beam is launched from the source, the tail of the beam is accelerated more than the head giving the beam a head-to-tail velocity gradient, or velocity “tilt”. This is done through the first five lattice periods. At this point, the beam is allowed to drift, letting the tail of the beam advance toward the head. This compresses the beam

I_{max}/I_o	$\epsilon_{max}^{rms}/\epsilon_o^{rms}$	a_{max}/a_o
1:1	0.9-1.1	1.0-1.2
1.5:1	0.9-1.1	0.9-1.0
3.7:1	1.3-1.6	1.6
7.4:1	≥ 2.8 -3.2	≥ 1.7 -2.0

Table 5.1: Results of the MBE-4 drift compression experiment. The data is taken from reference [17].

axially and raises the line charge density. Both mild and more aggressive compression schedules were run. The mild compression shots have a head-to-tail velocity tilt of about five percent, which increases the line charge density by roughly one half over fifteen lattice periods. The most aggressive shots have a head-to-tail velocity tilt of twelve to thirteen percent, increasing the line charge density by a factor of 7.4.

The mildly compressed shots show little or no emittance growth. The more aggressively compressed beams show significant emittance growth; growths up to a factor of three were seen. The transverse trace space plots for these shots show a spreading that resembles “bow ties” or “butterflies”. Near the edge of the beam, there is a significant spreading in transverse velocity. Table 5.1 gives the emittance growth versus compression factor and the emittance growth along the accelerator for the case with the highest compression. The trace space and beam profile for both high and low initial velocity gradients are shown in figure 5.1 and figure 5.2. The figures show sixteen plots of the trace space at increasing times along the beam, halfway between two quadrupole elements. The slope imposed by the alternating gradient motion was removed by subtracting it from x' . These experiments are discussed more fully in reference [16].

The MBE-4 lattice is a syncopated FODO lattice; one drift section is longer than the other. This was done to allow room for the diagnostics which were placed in the longer drift section. The following are the parameters describing the lattice and are the parameters used in the simulations.

Half period length = .2286 m

Quadrupole length = .1 m




Figure 5.1: Experimental results for the mildly compressed beam. The picture on the left is the beam profile (vertical axis is time). On the right are sixteen trace space plots at the times given. The trace space is fairly undistorted. These plots were produced by W.M. Fawley of LBL.




Figure 5.2: Experimental results for the aggressively compressed beam. The picture on the left is the beam profile (vertical axis is time). On the right are sixteen trace space plots at the times given. The trace space shows “wings” near the transverse edge of the beam. These plots were produced by W.M. Fawley of LBL.

Long drift length = .1857 m
 Short drift length = .0715 m
 Beam Aperture = 54 mm
 Electrode diameter = 40 mm
 Electrode end gap = 1 cm
 Voltages on quadrupoles = 0 V and -17100 V

5.2 WARP3d Model of MBE-4

An attempt was made to get close to the parameters that are typical to the MBE-4 experiments. There are some differences due to use of a matched beam in the envelope code (used for particle loading) so the results are not confounded by emittance growth and oscillations from mismatches. The following are the initial beam parameters used in WARP3d.

Current = 5 mA
 Energy = 180 keV
 Atomic Mass = 132.9 amu
 Charge State = 1
 Envelope size at center of long drift section = 7.52 mm
 Envelope angle at center of long drift section = 20.9 mrad
 Emittance = 16 mm-mrad

The code has all of the essential physics required for a semi-quantitative model of the MBE-4 experiment, including alternating gradient focusing, non-linear electrostatic self fields, and image charges. Some of the major details are examined.

The synchroated FODO lattice was included in WARP3d. The evenly spaced lattice was also used in the early runs, though differences in the lattice had little effect.

The quadrupole focusing can be modelled either ideally, with pure focusing fields, or realistically, with conductors. The ideal applied, external focusing fields are sharp

edged and have no fringe fields nor any higher order, non-linear multi-pole fields associated with real quadrupole elements, except the dodecapole field. The dodecapole field is the next higher field that has the same symmetry as the quadrupole field and is the dominant higher order field. On MBE-4, this field is 3.5% of the quadrupole field at the pole tip. The internal conductors give both fringe fields and higher order components of the fields, including the dodecapole field, implicitly as part of the field solution. See section 2.2.5 for details.

Two others effects that are part of the self-consistent field solutions are the energy effect and axial fields. The energy effect, the variation in the transverse plane of the electrostatic potential and hence particle axial velocity, is discussed in more detail in the chapter on modelling the ESQ injector. It is a smaller effect here than with the injector. The potential at the pole tip is about 5% of the beam axial kinetic energy. The axial fields are the result of both the inter-digital structure of the focusing quadrupoles and the relation between the ground potential of the pipe and the potentials of the focusing quadrupoles. They are thought to have little effect since they do nothing more than make the beam slightly mismatched from the variation in axial velocity. Also, in the experiment, the beams with lower compression ratio experience the axial fields (with only small transverse variation) and show no emittance growth.

In order to reduce the number of voltage sources for the electrostatic quadrupoles in MBE-4, one side of the quadrupoles were held at ground, while the others were charged to minus twice the necessary focusing voltage. This would be no different than charging at plus and minus the focusing voltage, except that the surrounding walls are held at ground. This was thought to have little effect on the beam behaviour. The major effect would be changing the axial fields which has minimal effect on the beam. Though the potential arrangement does change the higher order field components, the change is small. Overlays of the V_{00} and V_{44} components are shown in figure 5.3.

The effects of image charges in MBE-4 are included either implicitly or explicitly in WARP3d depending on the method of field solution. Images charges on the square pipe are included implicitly. The potentials on the square pipe walls are held at




Figure 5.3: Comparison of field components with anti-symmetric ($\pm V/2$) and nonanti-symmetric (0 and $-V$) quadrupole voltages. Solid is nonanti-symmetric voltages, and dashed is anti-symmetric. a) shows V_{00} , the potential on axis. b) shows V_{20} , and c) shows V_{44} . The rest of the components are nearly identical. The quadrupoles extend from $Z = -.13575$ to $-.03575$ and $Z = .03575$ to $.13575$.

ground, and no image charges are present on the walls. In the FFT field solver sine-sine transforms are done so the potential on the boundaries is zero. With the capacity matrix solver, used with round pipes and with simple quadrupole structures, the image charges are explicitly calculated from the capacity matrix. The image charges force the potential on the conductor surfaces to zero. With more complicated structures that require the SOR field solver, image charges are again included implicitly when the conductors are held at specified voltages.

Care was taken to be sure the beam was initialized correctly. When the simple ideal focusing model is used, the beam can be loaded in its entirety directly into the focusing channel since only transverse electric fields are applied; there are no applied fields in the z direction. In a self-consistent focusing lattice, though, there are fields in the z direction, from the energy effect and from the axial fields described above. When a beam is loaded directly into these fields, it rapidly becomes badly distorted. The axial fields cause bunching and debunching, leading to peaks and valleys in the line charge density. The energy effect causes a mismatch. Particles off axis have a differing axial velocity than those on axis and therefore receive differing focusing (or defocusing) impulses.

The problem was avoided by injecting the beam from a plane in the center of the long drift section. At that plane, there are essentially no external transverse fields (and so no energy effect) and no axial fields because of symmetry. The beam is injected with a constant axial velocity and a tilt is applied after the beam has been completely injected. To shape the ends of the beam, the number of particles injected in the end was varied so the parabolic falloff in λ was maintained. No effort was made to adjust the beam envelope in the ends—the same envelope used in the beam center was used in the ends. In order to compare like beams, the beams were injected in the same manner in runs with ideal focusing.

One problem with injecting the beam is obtaining correct axial self-fields. The head of the beam, which is injected first, starts blowing off after beam is injected behind it. Only the first ten centimeters out of 1.2 meters are affected, and the effect is an increase in axial velocity by less than 10%. The center of the beam is undisturbed. In addition, when a particle is injected, there is almost no space charge behind it and so it is decelerated by the space charge in front of it. This only lasts for at most one or two timesteps until enough beam is loaded behind the particle to balance the space charge in front. Depending on where the particles is initially placed within the injection region, it will have differing amounts of space charge behind it. Particles will thus receive differing amounts of deceleration, causing an axial thermal spread. The net drop in axial velocity is typically less than .1% of the axial velocity and the thermal spread less than .01% of the axial velocity. Both of these space charge effects are small and can be tolerated. If it were necessary to correct further, an artificial half-beam upstream of the injection region could be introduced by appropriate loading of the charge-density array.

Some effort was made to have the diagnostics of WARP3d match those of MBE-4. The WARP3d and MBE-4 diagnostics are most naturally computed in different frames of reference. The diagnostics of WARP3d are calculated as a function of z (axial coordinate) at a single time, while on MBE-4 they are measured as a function of time at a single z . The particles used in WARP3d's calculation of moments and in

the trace space plots are taken from within a finite range of z values. In MBE-4, the diagnostics are taken at a fixed location in the lab frame and can only be taken within a finite range of time. The best compromise in the code is to take fairly narrow ranges of z and to do the calculations frequently. The ranges have to be narrow enough to avoid longitudinal effects, such as rotation of the $x - x'$ (or $y - y'$) trace space ellipse from the alternating gradient focusing, but wide enough to contain enough particles for good statistics. The range that is usually used is 1 cm. The particle position and velocity are extrapolated in time to the centers of the range in z to further reduce any longitudinal effects. From this, data at diagnostic stations in the laboratory frame are saved to build up a profile versus time. Machinery for calculating moments at locations in the beam frame and the laboratory frame are described in section 2.1.4. The particle data that is extrapolated for the moments calculation can also be used in particle plots.

Another detail is the display of trace space plots. From the way the data are taken on MBE-4, the slant of the ellipse from the alternating gradient focusing is automatically removed. A parallelogram is laid over the ellipse of the beam and two-slit measurements made only within it. In WARP3d, the slope is calculated from the particle moments. In x , the slope, s_x is given by

$$s_x = \frac{\tilde{X}'}{\tilde{X}} \text{sign}(X\bar{X}' - \bar{X}\tilde{X}'). \quad (5.1)$$

The “sign” function returns the sign of its argument, the $\tilde{}$ are root mean square values, and the $\bar{}$ are average values. The data plotted on the vertical axis is then $x' - s_x x$ for the x trace space plots, and $y' - s_y y$ for the y trace space plots.

5.3 WARP3d Results

In the simulation of the drift compression experiments, there was no attempt to make a one-to-one comparison between the experiment and the simulations. To do this, one would need to include the accelerating fields which supply the axial tilt which can not be done cleanly in WARP. (The algorithm used to apply the accelerating fields




Figure 5.4: Time history of the line charge at mid-pulse of an ideal compressed beam. The solid line is for the most aggressively compressed beam.

assumes the the field is constant in time. As the fields needed to produce the axial tilt are time dependent, the application of the field may not be accurate.) One other major problem is the lack of data from the experiment. Most of the data taken was in one transverse plane. Little data was taken in the other transverse plane. This makes it difficult to obtain a beam distribution and any offsets from the axis in that plane.

We started the simulations with an idealized case. Semi-Gaussian beams (uniform in space, Gaussian in velocity) that were thin and well separated from the square pipe walls, with varying linear axial velocity gradients, were sent through a purely quadrupole, hard edged, focusing channel. The walls were 60 mm from the beam center. The simulation with 15% axial velocity gradient resulted in a line charge growth of nearly a factor of 5. The most aggressive compression on MBE-4 gave a line charge growth factor of 7.4 by comparison. Plots of the line charge density at the center of the beam and initial and final distributions for runs with 5%, 10% and 15% initial velocity tilts are shown in figures 5.4 through 5.8.

The runs showed only slight emittance growth, on the order of 10 - 20% for the most aggressive compression. This was due mostly to a numerical artifact associated with an insufficient number of particles. Increasing the number of particles reduced




Figure 5.5: Initial distribution before compression. The tilt shown in the $z - V_z$ plot is the “aggressive” tilt. The curl at the head of the beam is an artifact of the loading via injection.




Figure 5.6: Final distribution after compression with 5% velocity tilt.

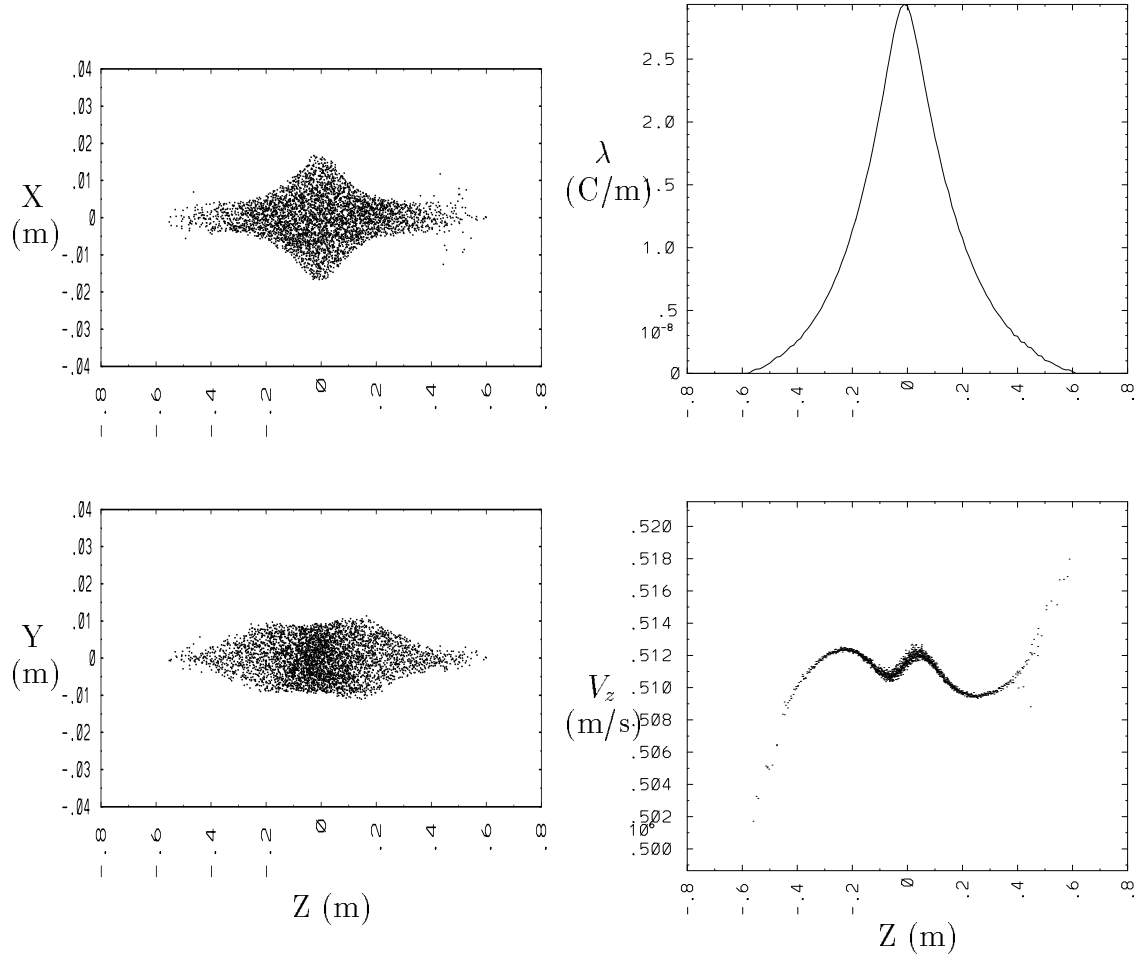


Figure 5.7: Final distribution after compression with 10% velocity tilt.




Figure 5.8: Final distribution after compression with 15% velocity tilt.




Figure 5.9: Time history of ϵ_x and $2X_{\text{rms}}$ of an ideal compressed beam. The mid-pulse emittance curves show the result for two different numbers of simulation particles. The top dashed curve is with 73574 particles and the bottom solid curve is with 294332 particles. The vertical axis is in normalized units.

the amount of growth. The x emittance of the center of the beam versus time for runs with differing numbers of particles is shown in figure 5.9. The trace space shows no distortion. Figure 5.10 shows the trace space for the most aggressively compressed beam. The behavior in the y direction is similar.

The growth in line charge was smaller since the axial fields are larger for a beam in a large-bore metal pipe than for a beam surrounded by quadrupole conductors. The axial fields can be written as

$$E_z \approx -g \frac{d\lambda}{dz} \quad (5.2)$$

where g is a geometrical factor and is given by $\ln \frac{b}{a}$. Here, a and b are the beam radius and the pipe radius respectively. Thus, for a larger metal pipe, g is larger and so the axial field is larger (the images in the pipe walls are so far away that they do little to counteract the self-field potentials). This retards compression, and so the beam turns around earlier and starts expanding. Figure 5.11 compares the axial potentials of a beam in various pipes. The highest curve is with the walls at 60 mm. The axial electric field, the gradient of the potential, is therefore the largest at the ends of the




Figure 5.10: $x - x'$ space of aggressively compressed ideal beam at four times in the middle of the drift section of the fifteenth lattice period. No distortion is evident. The times are relative to the time at which the center of the beam crosses the diagnostic location. The same AG slope was subtracted in all four plots. The figures at $\pm 2.15 \mu s$ show the change in the slope from the effect of the velocity gradient.




Figure 5.11: Self-potential of a beam in MBE-4 on axis. The double dashed line is with square walls at 60 mm and therefore has the highest value. The single dashed line is with square walls at 27 mm and therefore has the lowest value. The dashed line is with quadrupole rods only (pole tips at 27 mm). The solid line is with the full inter-digital structure including endplates. The last two also have surrounding walls at 60 mm. The peaks occur in the drift section where the surrounding wall is the only nearby conductor (as there are no quadrupole structures there).

beam.

Various perturbations were then added onto the semi-Gaussian distribution to examine possible causes of emittance growth. These included a ripple in v_z , to model possible high frequency errors in the accelerating fields on MBE-4, and mismatches, to model possible errors in the MBE-4 source. These two perturbations did result in additional emittance growth. A high frequency perturbation, with wavelength roughly one tenth the beam length, gave only slight emittance increase. A lower frequency perturbation, with wavelength equal to one half of the beam length, resulted in more growth, nearly 50%. Since this growth is less than what was seen in MBE-4, and more importantly, because the trace space plots did not show the distinctive “bow tie” shape, the velocity perturbation was thought not to be a major contributor to the emittance growth seen in MBE-4.

The beams were mis-matched by computing a mis-matched beam envelope solution and loading the beam with it. (The initial semi-major and minor radii of the envelope were changed to give the mis-matched solution.) This mismatch of the beam leads to additional emittance growth with the severity determining the amount of growth. A slight mis-match, changing the initial semi-radii of the envelope by less than 10%, results in no emittance growth. A more severe mis-match, changing the semi-radii by 50%, increased the emittance by 50%. Also, there was growth for both the mildly and aggressively compressed beams. These results led us to believe that the mis-match is not the cause of the emittance growth in the experiment.

Beams with a hollow distribution were also examined. Diagnostics from MBE-4 showed initial beams with non-uniform transverse distributions. Hollow beams were tried as an approximation. The transverse distribution that was used is given in equation 2.30. In the time scale of a plasma oscillation, the hollow beam collapses to a more uniform distribution. Most of the emittance growth, an increase by about 30%, occurs in this time. During the rest of the run, the emittance does not grow any more than with the initially uniform beams. The trace space plots do show a distinctive shape. The ends of the ellipse are bent in a counter-clockwise direction. The particles near the edge of the beam experience stronger space charge forces since more of the beam is away from the center. The focusing is lessened and the particles begin to lag behind in the rotation of the trace space ellipse. This behaviour is also seen in the experiment when beams are hollow. See figure 2 of reference [17].

5.3.1 Dodecapole and Image Fields

The higher order, nonlinear, dodecapole field was added to the externally applied focusing field. Beams that underwent mild compression showed little emittance growth in the presence of the dodecapole field. Beams that underwent more aggressive compression, though, did show moderate emittance growth. As a beam compresses, it expands transversely, maintaining a nearly constant charge density. Most of the emittance growth in these runs occurred near the end of the run when the beam radius




Figure 5.12: Time history of the ϵ_x and $2X_{\text{rms}}$ at the center of an aggressively compressed beam with an applied dodecapole field. The peaks near peak compression are a result of the dodecapole field. The peaks don't form until the beam radius expands to 18 mm.

expanded to more than 18 or 19 mm. Less aggressively compressed beams that started with a bigger radius also showed significant emittance growth as the beam expanded to a radius of 18 mm. At this point, the dodecapole electric field, which increases as r^5 , begins to affect the beam.

The increase in emittance observed for the most aggressively compressed beam was a factor of one third times the initial emittance. The history of the emittance is shown in figure 5.12. The trace space plots show evidence of the distinctive effect of the dodecapole. The trace space of the aggressively compressed beam in figure 5.13 can be compared to the trace space shown in figure B.4 of appendix B. The wings in the plot at 0 μs resemble the wings that appear in figure B.4. These appear since the beam has just exited a quadrupole which is focusing in y , so the dodecapole most affected the particles with large y .

The transverse beam profile formed a rectangle. This again is a result of the dodecapole field. The “tips” of the initial ellipse are pushed in from the additional focusing force of the dodecapole while the beam near the “corners” is pulled out from the dodecapole field which has the opposite sign. Figure 5.14 shows the profile at two




Figure 5.13: $x - x'$ space of an aggressively compressed beam with an applied do-decapole field at four times in the middle of the drift section of the fifteenth lattice period. The beam shows structures that resemble figure B.4.

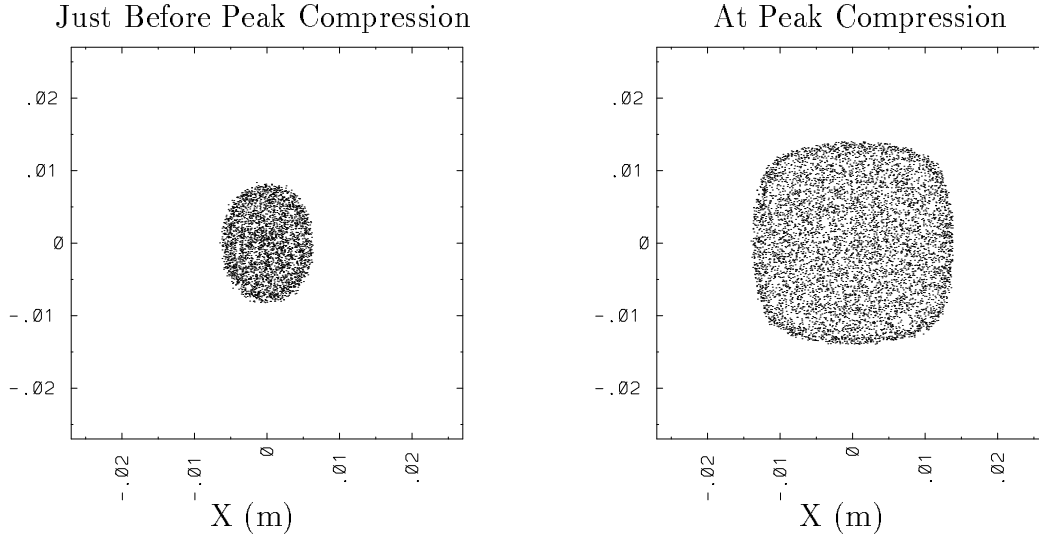


Figure 5.14: Transverse profile of the beam undergoing aggressive compression. The square shape is a result of the dodecapole. See text for more explanation.

locations along the accelerator. The shape is not the expected hexagon since when the dodecapole has the most effect, inside of a quadrupole, the beam is very elliptical. Only the parts of the beam near the ends of the major axis are affected.

The image charges and the geometric parameter, g , described above are next examined. The same initial runs as above were done but with the walls much closer in, 27 mm from the beam center, the location of the pole tips. This decreased the value of g . No dodecapole field was applied. The line charge did increase more, as expected. The line charge for the three compression schedules are shown in figure 5.15. The compression was slightly more than would be expected for MBE-4 since the square walls are on average closer to the center of the beam than the quadrupole conductors are. Also, in MBE-4 between quadrupole elements, the walls are much further away, further reducing the average g for the quadrupole conductors. In figure 5.11, the lowest curve, and hence the curve with the lowest gradient or axial electric field, is with the walls at 27 mm. Above it are two cases with quadrupole conductors present.

With gentle compression, little additional emittance growth was seen as the expansion of the beam did not bring it close to the wall. However, significant emittance




Figure 5.15: Line charge at mid-pulse of a compressing beam in a small pipe. The pipe walls are at 27 mm.

growth was seen with aggressive compression. The transverse expansion from axial compression puts the edge of the beam close to the walls. The images create a significant attractive force which distorts the trace space of the beam. Figure 5.16 compares the E_x field for a beam far away from the walls and for a beam close to the walls. The E_x field for a beam far from the walls is the linear self-field, with a drop-off outside the beam. With a wall close by, there is an additional expanding force from the attraction to the oppositely charged images on the wall; the field within the beam becomes non-linear. This non-linear force creates an essing in the trace space, which can be seen in figure 5.17. The figure also shows the resulting emittance history. The emittance growth seen in this run is higher than that which would be seen from images on quadrupole rod surfaces, since on average the square pipe will be closer to the beam than rods with pole tips at the same location.

Runs combining the dodecapole field with a small surrounding pipe resulted in emittance growth similar to that in the runs using either one separately. Figure 5.18 shows the emittance history and the resulting trace space near peak compression. With the increase in line charge at peak compression as compared to the case with the walls farther away, the beam radius also increases. This should lead to a stronger




Figure 5.16: Comparison of the self-potential of a beam close and far from the wall. The solid line is with the walls at 60 mm and is linear. The dashed line, with the walls at 27 mm, is nonlinear as a result of the image fields.

effect from the dodecapole. But the emittance does not increase more; the two effects seem to cancel each other. This result supports the conclusion of reference [18]. The image forces felt by the beam in this case, though, are not the same as felt by a beam within a quadrupole structure. A square pipe models the images of quadrupole conductors well when the beam is far from the conductor, but not as well when the beam gets closer. In the next section, the image forces are modeled more accurately.

5.3.2 Inclusion of Internal Conductors

To account correctly for the image charges on the quadrupole rods, the capacity matrix field solver was used to model the conductor geometry. The conductors were held at ground so their only effects would be images. The beam was focused with ideal, hard edged fields. As expected, the line charge does not increase as much as when the walls are at 27 mm. With the most aggressive tilt, the maximum line charge is $.056 \mu\text{C}/\text{m}$, as compared to $.066$ when the walls are at 27 mm. Also as expected, less emittance growth was observed than the case with walls at 27 mm. The emittance still increased by a factor of nearly two, though, and displayed the same essing that




Figure 5.17: $x - x'$ space and mid-pulse ϵ_x time history of a beam in a small pipe. The trace space plot is of the center of the beam at the center of the long drift section.

was seen in the square pipe.

With the application of the dodecapole, nearly all of the emittance growth went away. This further supports the findings of reference [18]. The resulting emittance histories are shown in figure 5.19. The trace space is still distorted somewhat, though, showing evidence of the effect of the dodecapole.

The internal conductors were then set to the focusing potential to supply the focusing self-consistently. This gives the dodecapole field as well as other high order multi-pole fields, fringing fields, images and the energy effect. The fields produced by the conductors were decomposed (with no beam present) to give the multipole fields. The dodecapole field in the center of a quadrupole is 3.5% of the quadrupole field at the pole tip, matching the MBE-4 value. Figure 5.20 shows the dodecapole field along one lattice period.

With aggressive compression, the emittance grows by a factor of nearly three by the time of peak compression. The beam was slightly mismatched and at peak compression had a larger rms x than the previous run, putting the beam edge closer to the wall. The effect of the dodecapole is more prevalent in the trace space— see figure 5.21. The cancellation between the dodecapole and image forces seems to break




Figure 5.18: $x - x'$ space and mid-pulse ϵ_x time history of a beam in a small pipe with dodecapole. The trace space plot is of the center of the beam at the center of the long drift section.

down as the beam get very close to the conductor and when other non-linear fields are present. The dodecapole has peaks at the end of the quadrupoles as shown in figure 5.20 but these peaks will not affect the beam much since, from the alternating gradient focusing, the beam is smaller at the ends of the quadrupole.

The beam was slightly mismatched because of the difference in the focusing fields between the ideal fields and the fields produced from the self-consistent quadrupole conductors. The fields in the centers of the quadrupoles are not the same due to error in the calculation of the field. See section 2.2.5 describing the capacity matrix field solver for more details. Also, the integral over one quadrupole (one half lattice period) of the focusing strength is not the same as the length of the hard edged quadrupole times its (constant) focusing strength. Furthermore, with the fringe fields, the focusing fields die off gradually away from the quadrupole. So, a beam loaded with an envelope solution obtained from matching to an ideal hard edged focusing lattice will not be matched to the self-consistent lattice. The mismatch was minimized when a matched envelope was found by scaling the focusing field so that the integrals over one quadrupole of the focusing strengths were the same.

The last step is to include the full inter-digital structure. This was done with the




Figure 5.19: Mid-pulse ϵ_x time histories with images on quadrupoles, with and without dodecapoles.

SOR field solver described in Section 2.2.5. The SOR solver can be quite expensive when used for time-dependent simulations. The run with aggressive compression took nearly two hours of CRAY C-90 CPU time as compared to about twenty minutes for capacity matrix simulations. The much richer physics of the inter-digital structure can only be captured by a first-principles model such as the one used in these runs.

The beam reaches a higher compression than in the above runs even though the image charges are the same. As shown in figure 5.22, λ reaches a maximum of nearly $.076 \mu\text{C}/m$, as compared to $.066 \mu\text{C}/m$ above. The peak occurs when the beam is in the short drift section. That point is in a trough in the axial potential from the inter-digital structure (location $z = 0$ in V_{00} of figure 5.3). The beam behind it is accelerating, while the beam ahead is decelerating, causing the beam to bunch. This creates a higher peak in the line charge. The two surrounding lower peaks in figure 5.22 are at the same spot in the previous and next lattice periods.

As mentioned above, the beam is mismatched and grows bigger than it would have otherwise. The mismatch is seen in the oscillation of the peaks of $2X_{\text{rms}}$ as shown in figure 5.23. The highest peak goes beyond 27 mm, the quadrupole aperture. As only




Figure 5.20: Dodecapole field with quadrupole rods. At $z = 0$ is a defocusing quadrupole and at the ends are focusing quadrupoles.




Figure 5.21: ϵ_x and $2X_{\text{rms}}$ at mid-pulse of a beam in self-consistent quadrupole rods.




Figure 5.22: Mid-pulse line charge time history of an aggressively compressed beam in the inter-digital structure.

a few particles are lost there the beam size did not go beyond 27 mm, and so $2X_{\text{rms}}$ is artificially high and the beam hollow.

Much emittance growth was observed in the simulation. This is due to two effects. The beam is mismatched and is large near peak compression, increasing the effect of the dodecapole and other non-linear components. The second effect are the other non-linear fields, specifically the components that are fourth order in radius. These cause curvature of the trace space and change the transverse shape of the beam, altering the image charges and altering how the dodecapole component effects the beam. The trace space plots for both x and y at the center of the long drift section are shown in figures 5.24 and 5.25. Note that the curvature is not strongly evident in those plots since it averages out over the course of a lattice period. At other points in the period, mainly within the quadrupoles, the curvature is evident.

The $y - y'$ trace space at the center of the long drift section clearly shows the effect of the fourth order potentials (as well as the dodecapole). Long tails result in some places and a bow tie shape near the center of the beam can be seen. The long tails appear only in y since the beam has just come from a focusing quadrupole in the y plane and is bigger in y than in x .




Figure 5.23: ϵ_x and $2X_{\text{rms}}$ at mid-pulse of a beam in in the inter-digital structure.

The bow tie appears to be a result of the dodecapole fields. Figure 5.26 shows the transverse profile and the $x - x'$ trace space near the center of the beam at the middle of the long drift section. The “+” and “-” symbols are the particles which have their shifted x' above and below zero respectively. The profile shows that all of the particles in two of the wings of the bow tie (the wing at negative x and positive x' and the wing at positive x and negative x') are off the y axis. The dodecapole would have that effect. The same is seen in the 5.27 for y .

5.4 Conclusions

Runs with WARP3d using idealized beams with idealized quadrupole focusing show less compression than is seen on the experiment, MBE-4, but showed little emittance growth. Perturbations to the initial distribution functions and non-linear transverse initial distribution functions increased the emittance of the beam, but the growth tended to be independent of the amount of compression. The effects on the trace space plots were unlike the effects seen on MBE-4.

The dodecapole field by itself has an effect on the emittance, as do the image forces by their self. When the two were combined, the effects seem to cancel. This supports

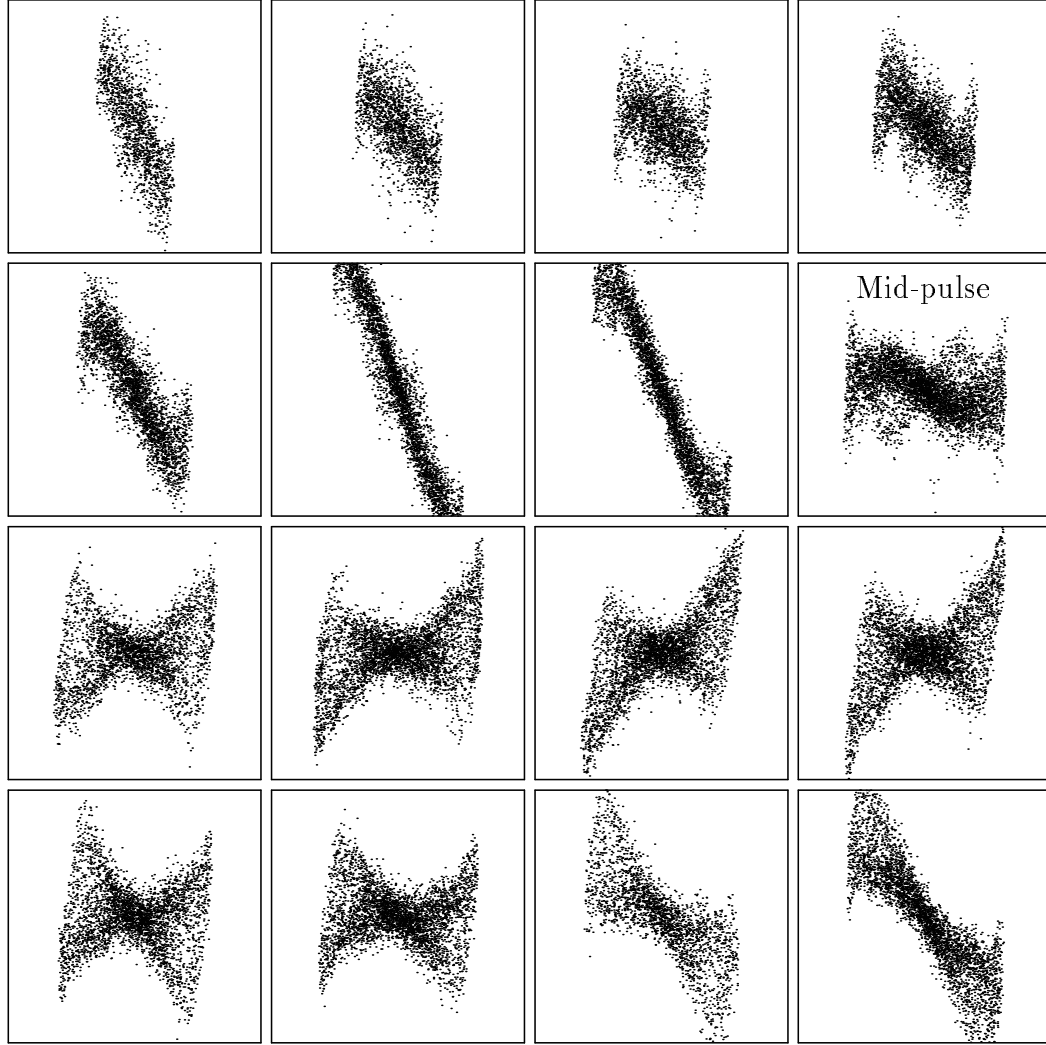


Figure 5.24: $x - x'$ trace space at sixteen place along the beam in the center of the long drift section. The frames are separated by $1.79 \mu s$ and frame number 8 is the mid-pulse of the beam. The minimum and maximum x (horizontal axis) are ± 27 mm and the minimum and maximum x' (vertical axis) are ± 5 mrad.

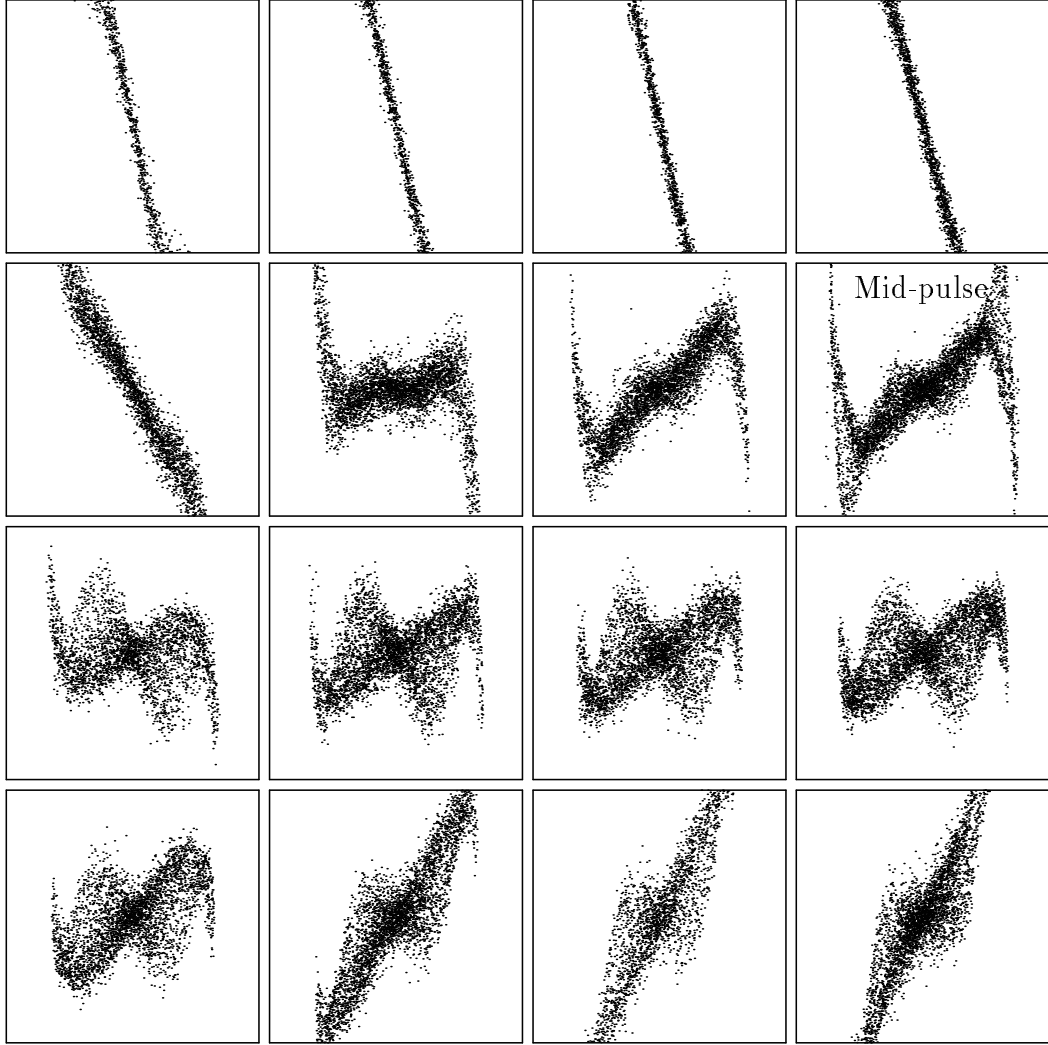


Figure 5.25: $y - y'$ trace space at sixteen place along the beam in the center of the long drift section. The frames are separated by $1.79 \mu s$ and frame number 8 is the mid-pulse of the beam. The minimum and maximum x (horizontal axis) are ± 27 mm and the minimum and maximum x' (vertical axis) are ± 5 mrad.




Figure 5.26: Beam x profile and trace space with symbols based on shift x' . This shows the effect of the dodecapole.




Figure 5.27: Beam y profile and trace space with symbols based on shift y' .

earlier results obtained using lower dimensional codes[18]. Both only affect the beam when the transverse expansion reaches nearly 19 mm. The cancellation appears to fail when the beam nearly touches the conductors; the effect of the dodecapole is evident in the emittance growth and in the distortions of the trace space.

The most significant simulation includes the full inter-digital structure. It includes dodecapole fields, images fields and octopole-like fields which result from the gaps between the quadrupole rod ends and the endplates. Much emittance growth is seen. The distortions on phase space show clearly the effects of the dodecapole fields as well as of the octopole-like fields.

Chapter 6

Electrostatic Quadrupole Injector (ESQ)

The injector for the ILSE (Induction Linac Systems Experiment) accelerator facility being planned at LBL will use electrostatic quadrupoles for both focusing and acceleration. An issue of concern is possible emittance degradation because the voltages applied are a significant fraction of the energy of the particles and there are significant nonlinear fields arising from the shape of the quadrupole structure. Simulations show that octopole and octopole-like fields ($\phi \sim r^4$) cause essentially all of the emittance growth. We believe that the effects can be reduced by adding correcting fields which can be created by altering the focusing structure and have pursued this.

6.1 Injector Description

The injector is required to supply 1 Amp of 2 MeV singly charged potassium ions at low normalized transverse emittance, less than $.5 \pi$ -mm-mrad. The source emits into a diode region, or gap, with a potential drop across and enters the first quadrupole. The potential drop across the gap gives the beam an initial kick. The size of the drop is critical to the design of the injector. This drop needs to be as large as possible so the beam energy is much larger than the focusing potential. However, physical




Figure 6.1: Diagram of an inter-digital quadrupole. a) is the transverse slice and b) is the side view. Also in b), the next quad is shown, giving the relative orientation of neighboring quadrupoles and what the voltages would be to supply a net accelerating voltage on axis. In the ESQ, there would be no gap between quadrupoles; they share a common plate.

constraints, such as voltage breakdown, put a limit on this field. This limit is between .5 MV and 1 MV. The beam then continues on through several more accelerating quadrupoles.

The quadrupoles have an “inter-digital” structure. See figure 6.1. Each has a plate on each end and four rods. With the origin along the center of the beam line, the two rods on the x axis are attached to one plate and the two on the y axis are attached to the other plate. The rod lengths are less than the distance between the end plates so there is a gap between the end of the rod and the plate to which the rod is not attached. When the plates are charged to different voltages, the rods produce quadrupole fields and a net change in potential along the axis. The effective length of the quadrupole field is proportional to the overlap of the rods.

To produce a net acceleration on axis, the high voltage end of each quadrupole is attached to the low voltage end of the previous. For example, if the first quadrupole

has the voltages 1.5 MV and 1.3 MV applied to the end plates, the next quadrupole would have the voltages 1.3 MV and 1.1 MV applied to its end plates. Each quadrupole would then have .2 MV across the rods to produce the quadrupole fields and the net potential drop across the two quadrupoles would be .4 MV. To produce alternating gradient focusing, the rods attached to the plate common between the two quadrupoles would be on the same axis.

With the high currents needed from the injector, large focusing fields are needed to confine the beam. The particles have low energy in the injector so the focusing fields are a large fraction of the particle energy. This results in what is called the “energy effect”. Particles that are off axis will be closer to one rod than the others, so their potential energy will then be significantly closer to the potential of that rod than to that of the others. If, for example, a particle is closer to the high voltage rod, it will be in a region with higher potential, and will therefore have a lower axial velocity than a particle on axis. This difference in velocity will be a large fraction of the total velocity, leading to beam distortion. If the distortion is uncorrected, it can cause emittance degradation. Note that this effect is present in all electrostatic quadrupoles, but the focusing fields are generally much smaller than the particle energy so the distortion is very small. The focusing structure produces fields that are not pure quadrupole. A variety of multipoles arise and are written as in equation 1.4.

6.2 Description of Computer Runs

The region of simulation starts within the diode. WARP3D currently does not have the capability of space-charge limited injection so the beam is injected with constant current. The point of injection is placed at a plane that is an equipotential. This is also done to facilitate boundary conditions; since WARP3D works in Cartesian coordinates, the boundary is given by a uniform potential on that plane. The simulations extend past the end of the injector to give a proper boundary condition on that end.

The runs are quasi time-dependent as only a steady-state solution is desired. The

quasi time-dependent method is the same as the fully time-dependent method except that the time step over which the field solves are done is longer than the time step over which the particles are advanced. At the start of the run, an initial bunch of particles is loaded into the injector and the fully self-consistent fields are calculated. Generally, the initial distribution loaded is cylindrical. These particles as well as particles injected near the source flow through the previously calculated fields for some number of time steps. The fields are then recalculated with the new positions and new particles. The particles are then moved through the new fields. This cycle continues until convergence, which generally is obtained shortly after one transit time across the injector. The transient behavior of the particles is lost—only the steady-state flow is obtained. The method is used for efficiency.

The field solutions are found using the PSOR method with internal conductors and subgrid scale boundary placement. The fourfold symmetry option is used (see section 2.2.5). The size of the field grid cells is limited by the need for accuracy of the self-fields and the field grid must be large enough to include enough of the quadrupole conductors to get good representation of the focusing fields. To meet these requirements, on the order of 50 grid cells are needed in the transverse direction; the beam occupies about a third and at least half of the rod is inside the grid. There are between 40 and 50 grid cells in the axial direction in each quadrupole. The total number of axial grid cells depends on the number of quadrupoles being simulated. The number is typically between two hundred and six hundred grid cells. The physical size of the grid cell depends on the scale of the system. In figure 6.2 the field grid is shown laid on the conductors and beam.

The injected particle distribution is a semi-Gaussian with a circular transverse profile. Typically, the number of particles used ranges from 70,000 to 300,000 total, with several hundred injected at each time step. Fewer particles lead to numerical problems and more particles are not necessary. Comparisons of runs with 70,000 and 300,000 particles showed little difference, though the runs with 70,000 were somewhat noisier.

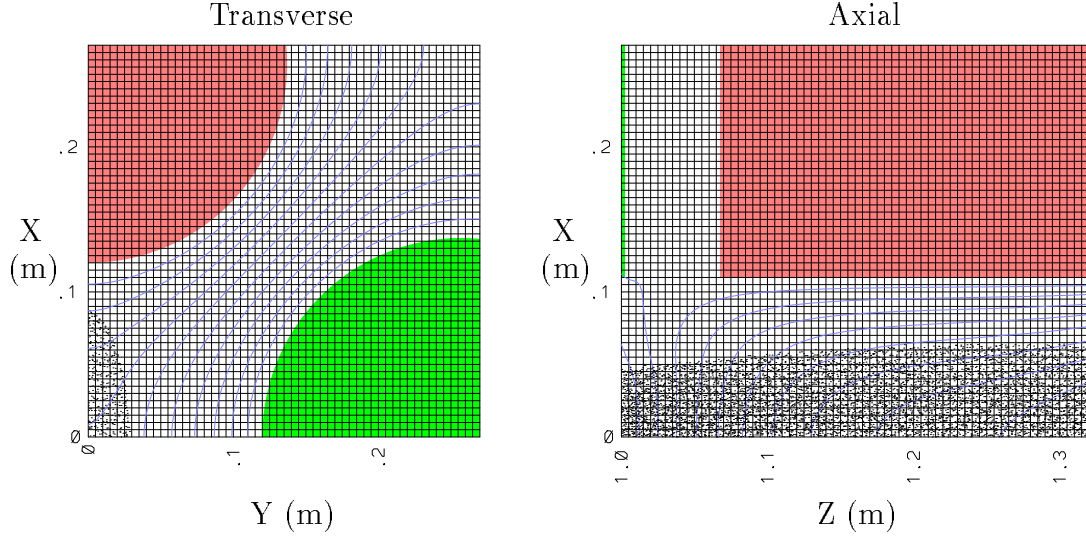


Figure 6.2: The field grid is shown here overlaying part of the system. The dots are the beam particles, the contour lines are lines of constant potential, and the gray areas are the conductors.

The transit time of typical runs is between 300 and 1000 time steps, depending on the time step size and number of quadrupoles. Another 100 to 200 steps are done to ensure convergence. The number of cycles, or field solves, was between five or ten. Most of the parameter searching runs were done with the bigger time step, 300 steps per transit time. Runs with a smaller time step were done occasionally to ensure that there were no numerical problems.

As a further check for numerical problems, fully time-dependent runs were made and compared to quasi-time-dependent runs. No difference was found. As check on convergence, longer runs were made. They were run over several transit times. Shortly after one transit time the beam reached a steady state and maintained it throughout the rest of the run.

The fields that arise from the conductors can be decomposed into the multipole components. This is done by first calculating the potential with no beam present. The V_{kl} of equation 1.4 are calculated by a least squares fit of the calculated potential to the truncated expansion. Typically, only terms out to $k = 6$ are calculated. This leaves ten terms and ten coefficients (V_{00} , V_{20} , V_{22} , V_{40} , V_{42} , V_{44} , V_{60} , V_{62} , V_{64} and V_{66}) in the expansion. The expansion is fit to the potential at a number of grid points

near the center of the beam. Typically 64 grid points are used in a square 8 points on a side with a corner of the square on the beam centerline.

All of the simulations were run on a CRAY C-90 YMP supercomputer. The CPU time used was as low as 3 minutes for runs with 70,000 particles and 5 field solutions over the course of the run for a total of 390 time steps. The bigger runs took up to 20 minutes of CPU time, with 280,000 particles, 10 field solutions, and 1560 time steps.

6.3 ILSE Scale Injector

The first simulations that were done were of the full scale ILSE injector with a .5 MV diode. The initial design of the injector was done by E. Henestroza of LBL with an envelope code. This code did not include the energy effect and the non-linear fields from the conductors. The main design criterion was to reduce the size of the envelope to minimize the effects of the non-linear fields. The parameters of the injector are given in table 6.1. The initial run with WARP3d showed a very large emittance growth. Figure 6.3 shows the normalized emittance along with the beam profiles and potential drop on axis. Note that the sharp rise in emittance in the y direction follows shortly after the widest part of the y envelope, and similarly in the x direction. The location $z = 0$ is the position of the actual source. In WARP3d, the beam is injected about half way between the true source and the first quadrupole, through a plane which is approximately an equipotential.

This emittance growth is not true emittance growth but results from the twisting or “essing” of phase space, which can be undone. The $y - y'$ phase space at several points along z is shown in figure 6.4. The first is at the point where y emittance is greatest. The second is near the middle of the third quadrupole element where the velocity from the AG focusing is small. The essing is clearly shown here. The third is near the end of the injector. Most of the essing vanishes, but the distortions have heated the beam somewhat. Long tails are also present.

The shape of the phase space, the essing, indicates action of fourth order (in

	ILSE injector	Scaled experiment
Beam current (Amps)	0.8	0.008
Initial beam radius (cm)	6.0	0.9
Quad voltages ^a (kV)	220.0	10.3
	260.0	12.2
	260.0	12.2
	255.0	11.97
	255.0	
	250.0	
Quad length ^b		
First quad	31.0	7.15
Others	47.0	11.155
Rod length (cm)		
First quad	25.5	5.89
		6.195 ^c
Others	39.5	9.7
Apertures ^d (cm)	13.0	3.25
	12.0	3.0
	11.0	2.75
	10.0	2.5
	9.0	
	8.0	
	8.0	
Plate width (cm)	1.0	.6

^aThe ILSE scale injector has either four, five, or six quads depending on the diode voltage. The scaled injector has four quads.

^bDistance between center of plates.

^cThe rods of the first quad are not the same length.

^dDistance from beam center to the rod. Rod radii are 8/7 of the aperture.

Table 6.1: Parameters for ILSE injector and scaled injector experiment. The ratio of the diode voltages is the same as the ratio of the quadrupole voltages.

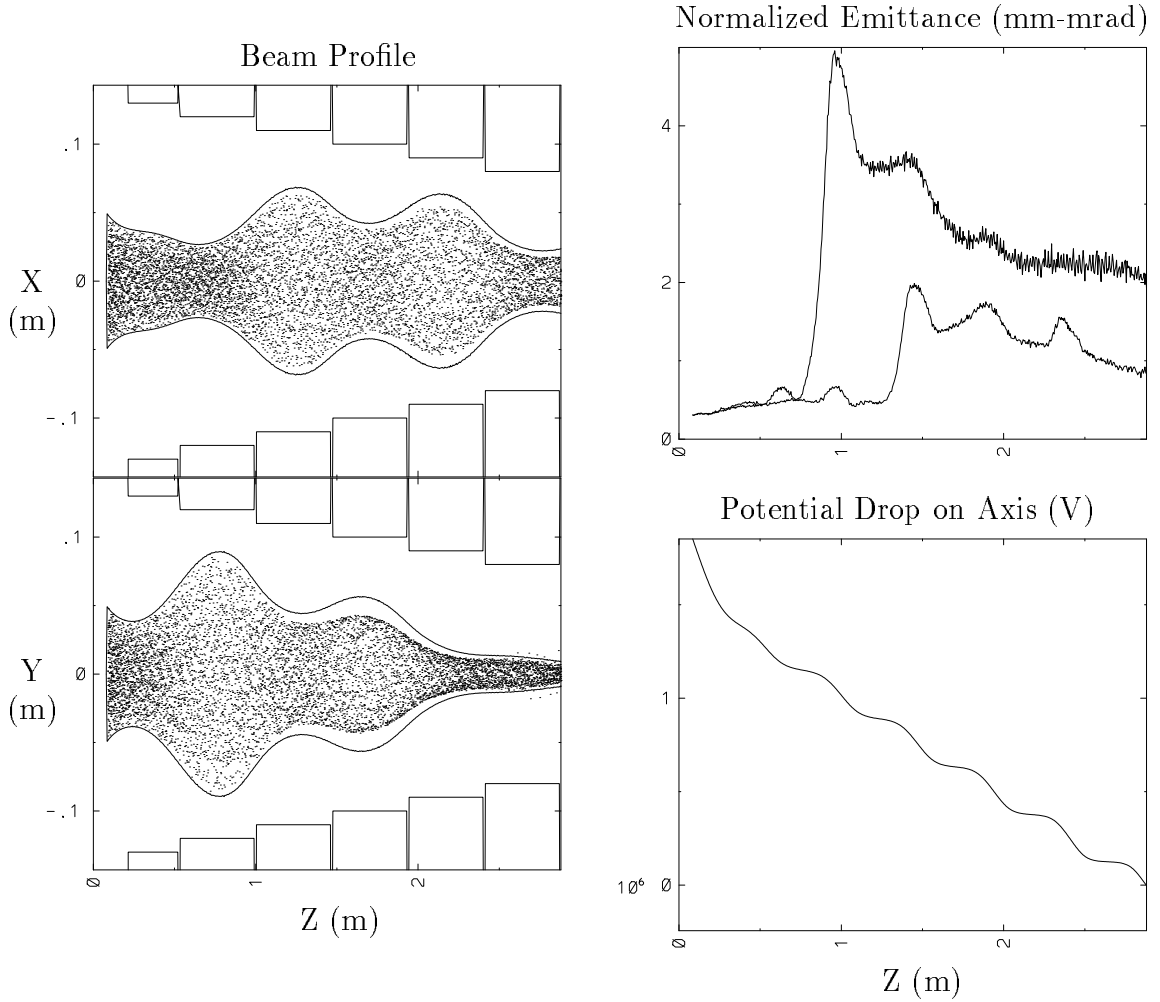


Figure 6.3: Initial ILSE injector results. The top curve in the emittance corresponds to the y profile. The curves surrounding the beam profiles are the RMS envelopes. The rectangles are the quadrupoles. Note that the peak in emittance follows shortly after the widest part of the profile. The gaps between the particles and the RMS envelopes result from bunching near the beam edge.

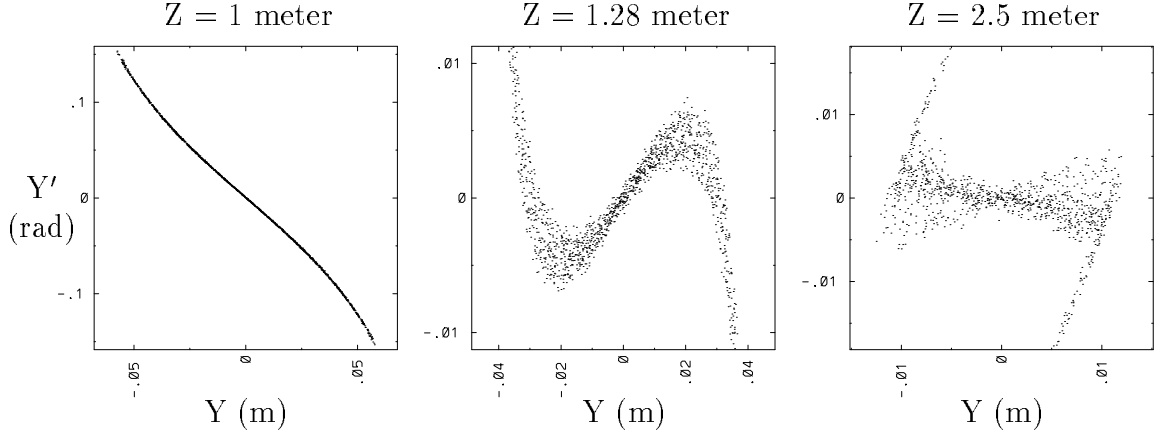


Figure 6.4: ILSE injector phase space at various locations along the injector. The peak in the Y emittance occurs at $Z = 1$ meter.

potential) fields, ϕ_{4m} . See Appendix B. The same runs as above were carried out, but cancelling out the various fourth order fields. The negative of the fourth order fields, obtained from multipole decomposition, was applied to the beam to cancel the fields. Cancelling the pure octopole field, ϕ_{44} , and the ϕ_{40} field had little effect on the emittance growth. Cancelling the ϕ_{42} field however, had a dramatic effect on the emittance; it was cut in half. Figure 6.5 shows the resulting emittance and phase space.

Plots of the fourth order field strengths are shown in figure 6.6. Both V_{40} and V_{44} change sign over one quadrupole element. Coupled with the fact that the energy changes little inside a quadrupole, the effect on the beam nearly vanishes with integration over one quadrupole element. The V_{42} , on the other hand, does not change sign over one quadrupole element. With integration over one quadrupole, its effect builds up. However, its sign does change from quadrupole to quadrupole, partially cancelling itself. The cancellation is not complete since the envelope and the energy of the beam change. The large effect in y of the ϕ_{42} field of the first quadrupole is partially cancelled by subsequent quadrupoles and the emittance decreases.

The remaining emittance growth is caused by the energy effect. In figure 6.7, $z - v_z$ phase space, the energy effect can be seen. It is the broadening in the flat portions, which are inside the quadrupoles. The effect is most evident in the second

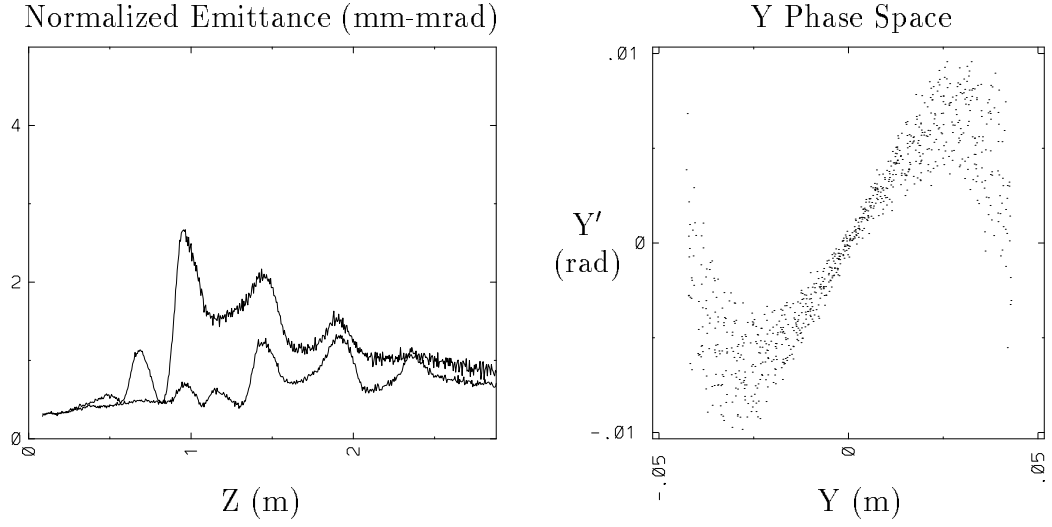


Figure 6.5: Injector results with V_{42} cancelled. The higher curve in the emittance is the Y emittance.




Figure 6.6: Fourth order fields in the ILSE injector.




Figure 6.7: Axial phase space in ILSE injector. The periodic broadening is from the energy effect.

quadrupole, between .5 and 1 meter. The beam energy is low and the beam profile is fairly wide, enhancing the spreading of v_z .

The energy effect was analyzed by S. Yu of LBL by studying the equation of motion of a particle in the fields. The Lorentz equation for v_x in the potential ϕ is given by

$$\frac{dv_x}{dz} = -\frac{e}{m} \frac{\partial \phi}{\partial x} \left(\frac{1}{v_z} \right), \quad (6.1)$$

where ϕ is as is in equation 1.4,

$$v_z = v_{z0} \sqrt{1 - \frac{e\phi}{T}} \quad (6.2)$$

is the velocity along the axis,

$$v_{z0} = \sqrt{\frac{2T}{m}} \quad (6.3)$$

is the axial velocity on axis, and T/e is the potential drop on axis. The $\frac{1}{v_z}$ results from the conversion from dt to dz . The expansion of equation 6.1 to second order in potential results in the linear equations of the motion of a particle in quadrupole fields. Expansion to fourth order in potential gives the perturbations from the non-linear fields. The resulting expression for the perturbation in v_x is

$$\frac{dv_{x1}}{dz} = \frac{v_{z0}}{2} \left\{ \left[\frac{1}{8} \left(\frac{V_{22}}{T_0/e} \right)^2 + \left(\frac{V_{44}}{T_0/e} \right) \right] \left[-\frac{\partial}{\partial x} \left(\frac{r}{R} \right)^4 \cos 4\theta \right] + \right.$$

$$\left[\frac{1}{2} \left(\frac{V_{22}}{T_0/e} \right) \left(\frac{V_{20}}{T_0/e} \right) + \left(\frac{V_{42}}{T_0/e} \right) \right] \left[-\frac{\partial}{\partial x} \left(\frac{r}{R} \right)^4 \cos 2\theta \right] + \quad (6.4)$$

$$\left[\frac{1}{8} \left(\frac{V_{22}}{T_0/e} \right)^2 + \frac{1}{4} \left(\frac{V_{20}}{T_0/e} \right)^2 + \left(\frac{V_{40}}{T_0/e} \right) \right] \left[-\frac{\partial}{\partial x} \left(\frac{r}{R} \right)^4 \right] \Bigg\}$$

The terms in $V_{2m}V_{2n}$ arise from the interaction of the terms $\frac{\partial \phi}{\partial x}$ and $\frac{1}{v_z}$ which represents the energy effect. It is the interaction of the focusing fields and the energy difference between off axis particles and the beam center. The terms in V_{4m} arise from $\frac{\partial \phi}{\partial x}$, the quadrupole structure.

The energy effect can be cancelled as the fourth order fields were, by applying the appropriate fields to the beam. The correcting fields should, on integration over one quadrupole, cancel the energy effect. Over one quadrupole, the term $V_{22}V_{20}$ changes sign; its integral nearly vanishes. The term $(V_{20})^2$ is small compared to $(V_{22})^2$ and so can be ignored. This leaves only $(V_{22})^2$ as significant. The correcting fields that are then needed are V_{44}^{cor} and V_{40}^{cor} .

Since $(V_{22})^2$ has the same behavior in the first and third lines of equation 6.4, V_{44}^{cor} and V_{40}^{cor} should have the same value. The sizes of V_{44}^{cor} and V_{40}^{cor} , which are scaled by V_{22} , are given for each quadrupole by ϵV_{22} , where

$$\epsilon = -\frac{1}{8} \int_{\text{quad}} \left(\frac{V_{22}}{T_0/e} \right)^2 dz / \int_{\text{quad}} \left(\frac{V_{22}}{T_0/e} \right) dz \quad (6.5)$$

The applied fields are $V_{40}^{\text{cor}}(r/R)^4$ and $V_{44}^{\text{cor}}(r/R)^4 \cos 4\theta$. They are scaled by V_{22} with the idea that they would be applied along with or from the quadrupole rods.

The correcting fields reduce the emittance growth by nearly a factor of two. The decrease in emittance is continuous along the injector after the first peak. Compare the emittance in figure 6.8 to the emittance with no applied corrections, figure 6.3. The initial peak from the ϕ_{42} fields is still present. The phase space in the third quadrupole is also shown in figure 6.8. The tails of the esses are shorter than in figure 6.4.

One more run was done where both corrections were done; the fourth order fields were cancelled and the energy effect was corrected. The plot of the emittance, figure




Figure 6.8: Injector results with energy effect correction. The higher curve in the emittance is the Y emittance.

6.9, shows that there is essentially no emittance growth and all of the emittance in the phase space has disappeared. We conclude that essentially all of the emittance growth resulted from the fourth order potentials and energy effect.

6.3.1 1 MV Diodes

The injector was also simulated with a diode energy of 1 MV. It is expected that this case will show less emittance growth than the .5 MV case. By increasing the energy before it enters the ESQ, the beam is stiffer and maintains a smaller envelope, reducing the effect of both the nonlinear fields and the energy effect. Also, with a higher beam energy, the energy effect is directly reduced. The results of the simulation of an optimized injector with 1 MV diode are shown in figure 6.10. The focusing voltages were optimized to reduce the size of the envelope. As expected, the growth is much less than with the .5 MV diode. Application of the correcting fields further reduced the emittance as in the .5 MV case. This result shows that if a 1 MV diode can be realized in the experiment, very little further work is needed to optimize the injector.

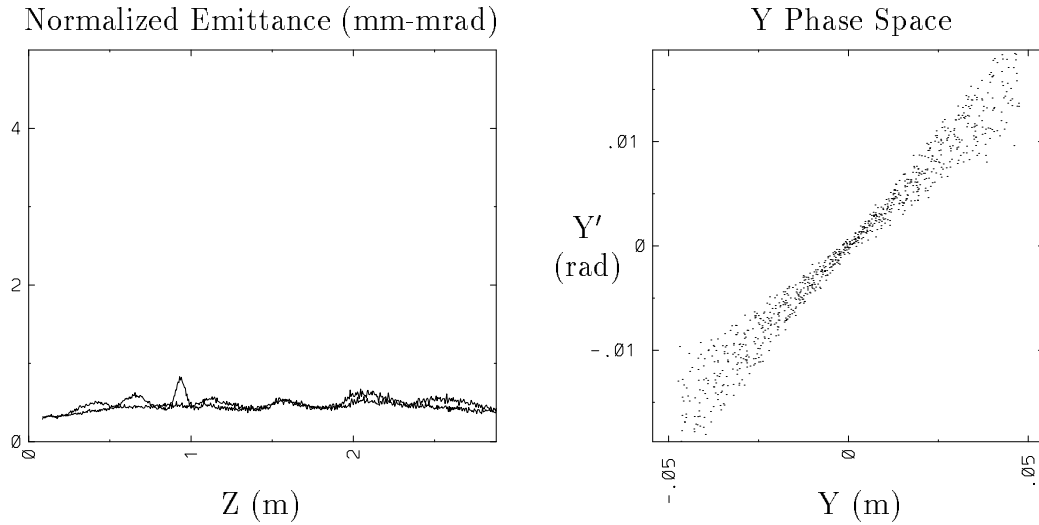


Figure 6.9: Injector results with both corrections. The emittance becomes flat and the essing of the phase space goes away.

6.4 Self-Consistent Corrections and Optimization

The goal of these simulations was to reduce the emittance from the injector to an acceptable level. We have shown that the growth is purely a result of external fields and that it is possible to reduce the emittance growth by altering those external fields. No beam manipulations need to be done. We then sought to do this in a self-consistent manner, altering the conductor geometry to alter the external fields appropriately. This was done for a scaled .5 MV diode injector to match the injector experiment performed at LBL.

The parameters of the experiment are given in table 6.1. The physical dimensions are scaled by a factor of four, and the voltages are scaled by a factor of roughly 21. In this experiment, the beam passes through a grid just before it enters the quadrupoles. The grid is an equipotential and is used as the injection plane in WARP3d. At the end of the injector, a plane at ground potential is used to represent the experimental diagnostic. (It is either a two-slit device to measure phase space, or a Faraday cup to measure current.)

The first step was to further optimize the envelope and to make it as small as

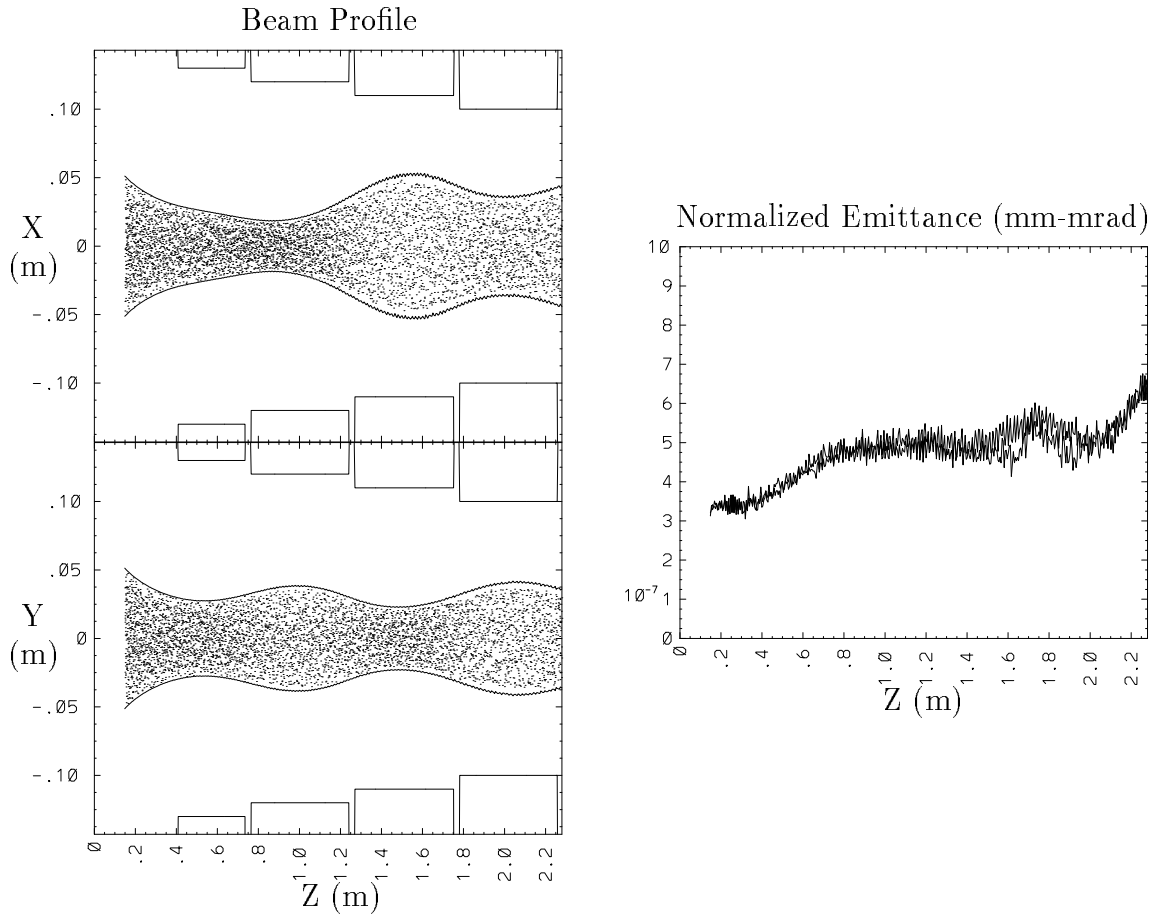


Figure 6.10: Injector results with a 1 MV diode. The emittance growth is much smaller than with the .5 MV diode. The quadrupole voltages are 176kV, 338 kV, 338 kV, and 255 kV.

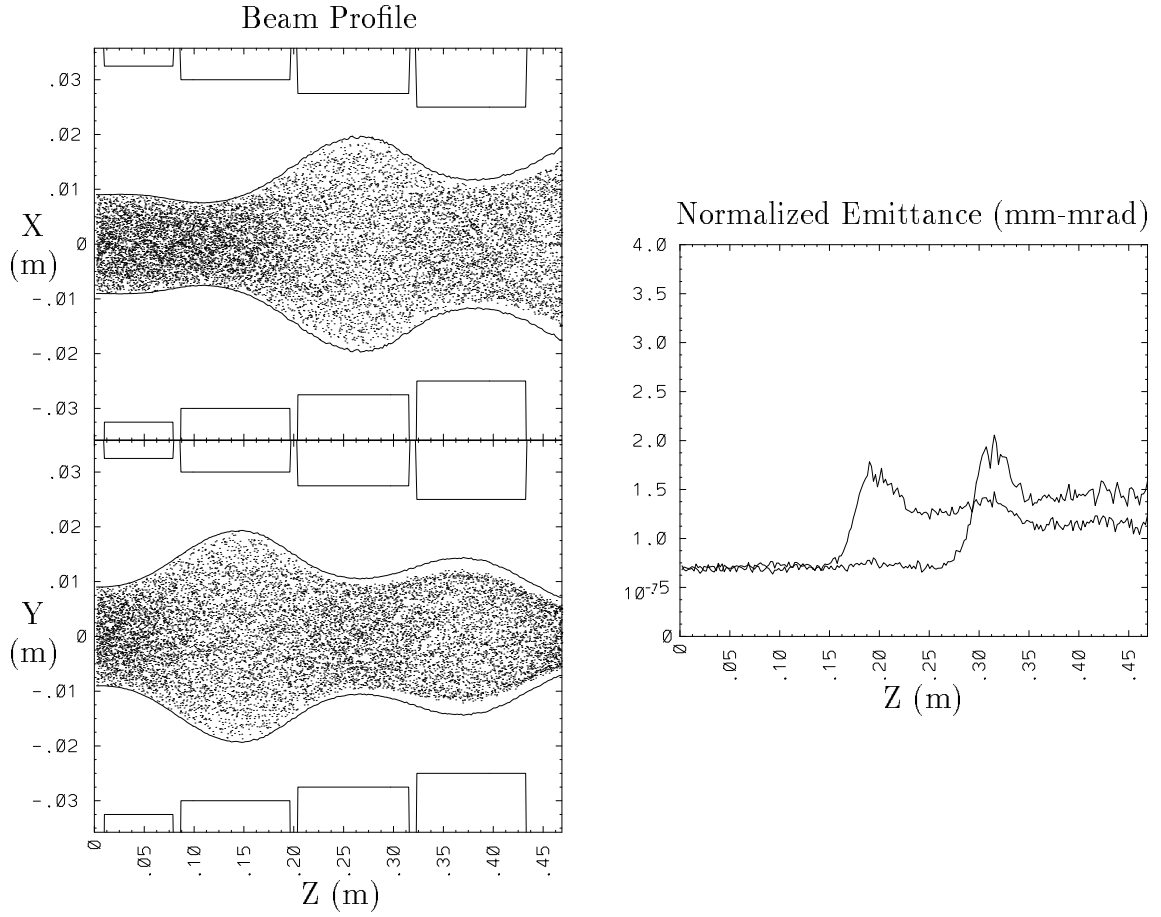


Figure 6.11: Scaled .5 MV injector with optimized voltages. The voltages are optimized to minimize the size of the envelope.

possible. That reduced the emittance by only a small amount as it was already somewhat optimized with the envelope code. The minimization of the envelope is limited. The best solution that could be obtained for the scaled injector with the .5 MV-equivalent diode is shown in figure 6.11. The quadrupole voltages and apertures are given in table 6.2. Though limited, this optimization reduced the emittance to a satisfactory level in the injectors with a 1 MV diode. In the scaled experiment plots, the location $z = 0$ is the edge of the first plate.

Fourth order fields can be altered by changing the structure of the quadrupole.

	Quad Voltages (kV)				Quad Offsets (mm)				Final ϵ_x (mm-mrad)	Final ϵ_y (mm-mrad)
	1	2	3	4	1	2	3	4		
a	10.3	12.2	12.2	11.97					.25	.17
b	12.36	13.42	14.03	9.87					.125	.15
c	12.36	13.42	12.81	9.87	-2.4	-4.8	3.6	-7.2	.1	.121

Table 6.2: Quadrupole voltages and offsets for .5 MV scaled injector. Row a is the voltages scaled from the ILSE injector. Row b are the optimized voltages. Row c are the optimized voltages and quadrupole offsets

It must be done in such a way that the integral of the quadrupole field over each element remain the same; the focusing and therefore the beam envelope will then remain optimized. The most effective structural change was found to be moving the focusing rods out and the defocusing rods in. The rods are moved away from the fattest parts of the beam. (The beam is at its largest inside the focusing quadrupoles.) Many runs were done in the search for the optimal changes. The final offsets obtained are given in table 6.2 and the results are shown in figure 6.12.

Some of the multipole fields from the quadrupole configuration with shifted rods are compared to the fields from the configuration with unshifted rods in figure 6.13. The field on axis, V_{00} , is lower with offsets since the defocusing rods are moved closer to the axis and they have the lower voltage of the quadrupole conductors. The quadrupole fields have differing shapes, but the integrals over each element are the same. The V_{42} field is shown since it is the one which has the most effect on the beam. The effect of the offsets on V_{42} is to increase the field on entrance to an element and to decrease it on exit. This reduces the effect of the field since it is cancelled more completely. At the exit of an element, the beam feels a weaker V_{42} field. The beam then accelerates between elements and then feels at the entrance of the next element a stronger V_{42} that is of opposite sign than of the previous element. Since the beam is stiffer from the acceleration, the effect of the field is weakened, so the stronger V_{42} cancels the V_{42} field from the previous quadrupole more completely.

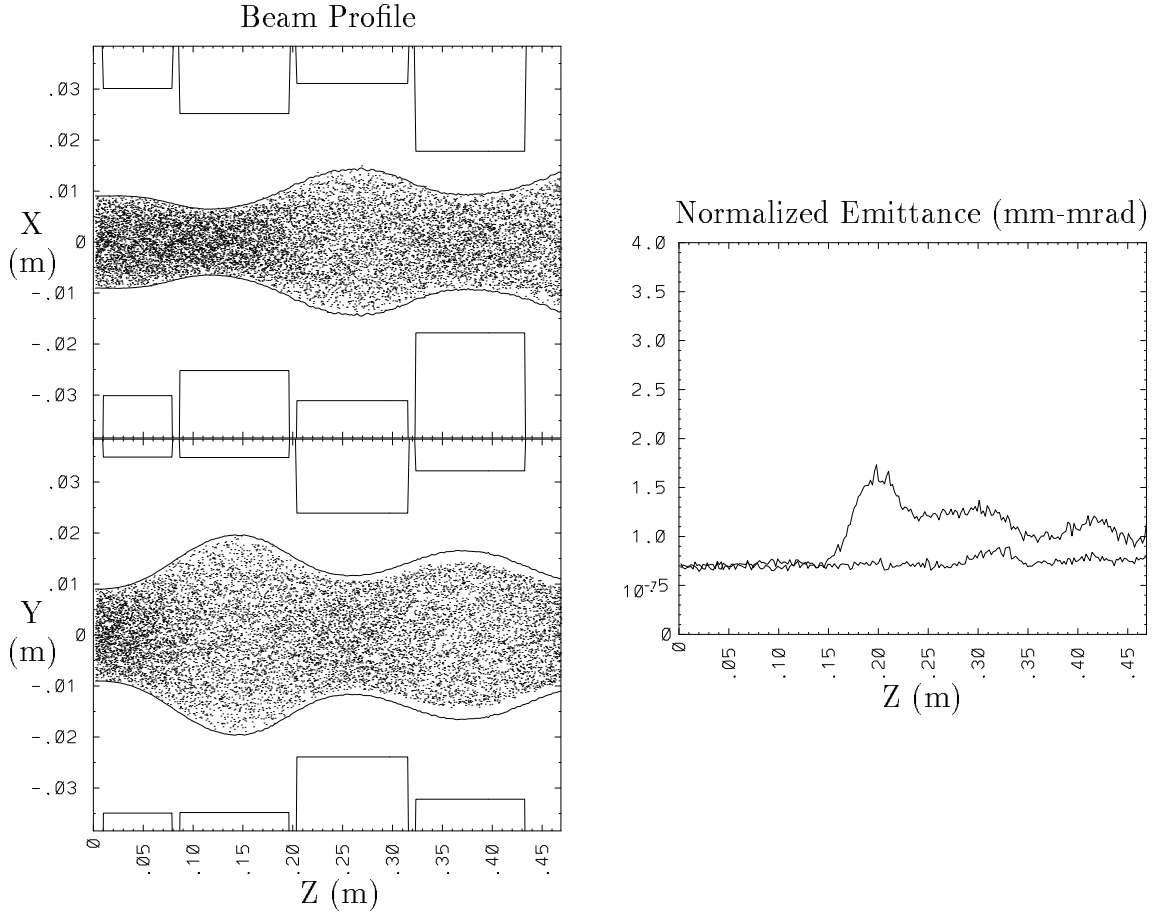


Figure 6.12: Scaled .5 MV injector with optimized voltages and rod offsets. The voltages are optimized to minimize the size of the envelope. The focusing rods are shifted in and the defocusing rods are shifted out. This moves the rods farther away from the beam, reducing the effect of non-linear fields, while maintaining the same quadrupole focusing field.




Figure 6.13: Comparison of multipole fields in the scaled injector with and without offsets. The dashed curves are with the offsets given in table 6.2. See text for more explanation.

6.5 Comparison With Experiment

As a proof of principle experiment, a scaled down version of the ILSE injector was built and tested at LBL. The experiment had the parameters as given in row a of table 6.1 and as used in the simulations in the previous section. Two series of experiments were performed; one a scaled version of the injector with the .5 MV diode, the other a scaled version of the injector with the 1 MV diode.

In the experiment, the two-slit emittance diagnostic did not have the same axial position for the measurements of the results in the x direction and in the y direction. The results in the x direction were taken with the diagnostic 1.1 inches past the end of the fourth quadrupole. The diagnostic was .88 inches past the end of the fourth quadrupole for the measurements in the y direction. The diagnostic was modeled in WARP3d by a plane that was held to ground potential.

6.5.1 500 kV Diode

For the .5 MV injector, a series of runs were done to understand the sensitivity of the design to changes in the initial beam size and quadrupole voltages. The results are shown in tables 6.3 and 6.4. In both cases, the diode voltage was 23.68 kV. The results show that the beam can be fairly sensitive to changes. A decrease in the quadrupole

V_{fac}	x (cm)	y (cm)	x' (mrad)	y' (mrad)	ϵ_x (mm-mrad)	ϵ_y (mm-mrad)
0.8	17.5	9.5	80.	-110.	.2	.16
0.9	15.0	7.0	80.	-85.	.16	.12
0.95	14.0	5.0	80.	-80.	.14	.12
1.0	12.5	4.8	80.	-70.	.13	.12
1.05	12.0	3.7	77.	-50.	.11	.13
1.1	11.0	3.0	75.	-40.	.10	.14
1.2	10.	2.0	80.	-20.	.09	.17

Table 6.3: Results of the .5 MV scaled injector with varying quadrupole voltages. All of the quadrupole voltages are varied by the factor V_{fac} .

a_{fac}	x	y	x'	y'	ϵ_x	ϵ_y
0.8	13.5	6.5	78.	-90.	.21	.11
0.9	13.5	5.5	79.	-80.	.15	.10
0.95	13.5	5.3	80.	-75.	.14	.11
1.0	12.5	4.8	80.	-70.	.13	.12
1.05	12.5	4.1	78.	-60.	.12	.12
1.1	12.0	3.7	78.	-57.	.10	.13
1.2	12.0	3.0	75.	-47.	.09	.17

Table 6.4: Results of the .5 MV scaled injector with varying initial beam radius. The initial beam radius is varied by the factor a_{fac} .

voltages or initial beam radius causes nearly a factor of over one and a half increase in the emittance in both directions. On the other hand, an increase in quadrupole voltage of initial beam radius tends to cause a small decrease in emittance. A decrease in the voltages will let the beam envelope grow bigger, increasing the effect of the non-linear terms. An increase in the voltage will tend to make the beam envelope smaller, decreasing the effects. A decrease in the initial beam radius will make the beam expand initially and it will overshoot the nominal case. The beam will thus be effected more by the non-linear fields. An increase in initial radius will let the beam shrink and undershoot the nominal case. The affect of the non-linear fields will be lessened. It is better to set the focusing voltages and initial beam radius too large.

On the experiment, data were taken with varying diode voltage. The series of runs was simulated. Table 6.5 gives the initial data. The final results are compared

diode voltage (kV)	current (mA)	ϵ (mm-mrad)	radius (mm)	angle (mrad)
19.	4.82	.07	9.0	0.0
25.	7.28	.07	9.0	0.0
30.	9.57	.07	9.0	0.0
35.	12.06	.07	9.0	0.0

Table 6.5: Initial data for the .5 MV scaled injector.

in figure 6.14. The simulation results agree well with the experiment.

6.5.2 1 MV Diode

The parameters for the 1 MV diode scaled experiment were obtained with WARP3d. (The same physical setup as in the .5 MV experiment was used.) The reduction of noise in the current measurement put a lower limit on the current, thereby imposing a lower limit on the diode voltage also. A value of 30 kV was chosen. To keep the same perveance, $I/V^{\frac{3}{2}}$ was kept the same in the scaling (I is the beam current, V is the diode voltage). This gave a current of 4 mA, which was large enough to overcome the noise in the experiment. To obtain this current in the experiment for the given diode energy, the beam from the source was scraped down to 8.5 mm at the entrance to the first quadrupole (at the grid).

With those as starting values, the quadrupole voltages were varied in WARP3d simulations to obtain the optimal, or minimal, envelope and emittance growth. Roughly fifty runs were made, each using minimal numbers of particles and field solutions and each requiring 3 minutes of C-90 CPU time. The quadrupole voltages that resulted are 7.392 kV, 9.126 kV, 9.43 kV, and 9.18 kV. Note that these values are not the same as the optimized voltages obtained for the unscaled 1 MV injector as given in figure 6.10 since the quadrupoles are not exactly scaled. The rest of the initial data is given in table 6.6

As in the 500 kV scaled case, data was taken for various diode voltages. The results are compared in figure 6.15. The two agree well except in the emittance where the errors are within the experimental uncertainty. Figures 6.16 through 6.19

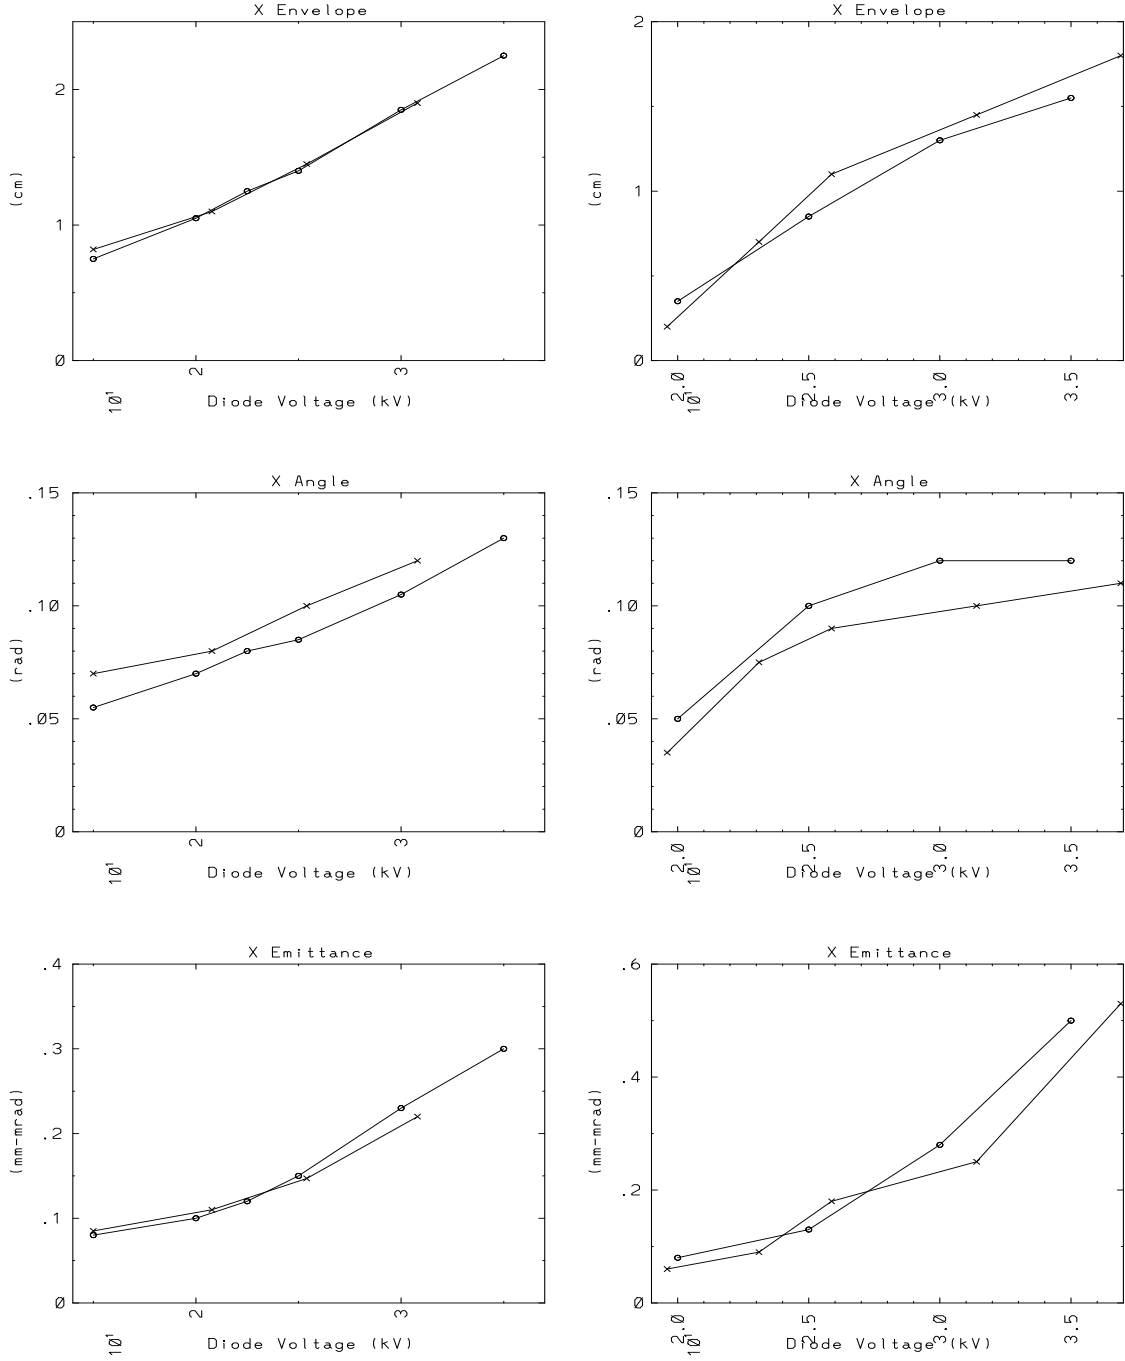


Figure 6.14: Comparison of WARP3d and experiment for the 500 kV scaled injector. The o's are the WARP3d results and the x's are the experimental results. The left column compares the x values for the set of voltages in row (a) of table 6.2. The right column is with the quads voltage at .9 times the row (a) voltages.

diode voltage (kV)	current (mA)	ϵ (mm-mrad)	radius (mm)	angle (mrad)
15.	1.3	.04	8.5	0.
20.	2.08	.045	8.5	0.
25.	2.92	.075	8.5	0.
30.	3.9	.1	8.5	0.
35.	4.88	1.2	8.5	0.

Table 6.6: Initial data for the .5 MV scaled injector.

compare the trace space of the two. The trace space data was taken on a rectangular grid that was slanted to match the alternating gradient slope. At each grid point is a horizontal line whose length is proportional to the charge density at that point.

6.6 Conclusions of ESQ Injector Study

Study of the ESQ injector has given the result that, in the regime of interest, nearly all of the emittance growth is the result of external fields only, and only fields which are third order in r (potentials which are fourth order in r). The result is encouraging since it affords the possibility of reducing or eliminating the emittance growth by only changing the external fields. The external fields can be changed fairly easily by adjusting the geometry of the conductors. Unfortunately, adjusting the conductors to exactly cancel the nonlinear fields is difficult, though the self-consistent simulations, including the conductors, show that much progress can be made in reducing the emittance via simple changes in the geometry. The largest reduction (by changing the geometry only) is obtained by moving the focusing rods farther from the beam while moving the defocusing rods closer to the beam.

The simulations also show that a higher voltage in the diode decreases the emittance growth. The beam has a higher energy in the quadrupoles so the energy effect is smaller and the envelope is smaller. Both effects lead directly to emittance reduction. However, experimental difficulties, such as voltage breakdown, put a limit on the diode voltage.

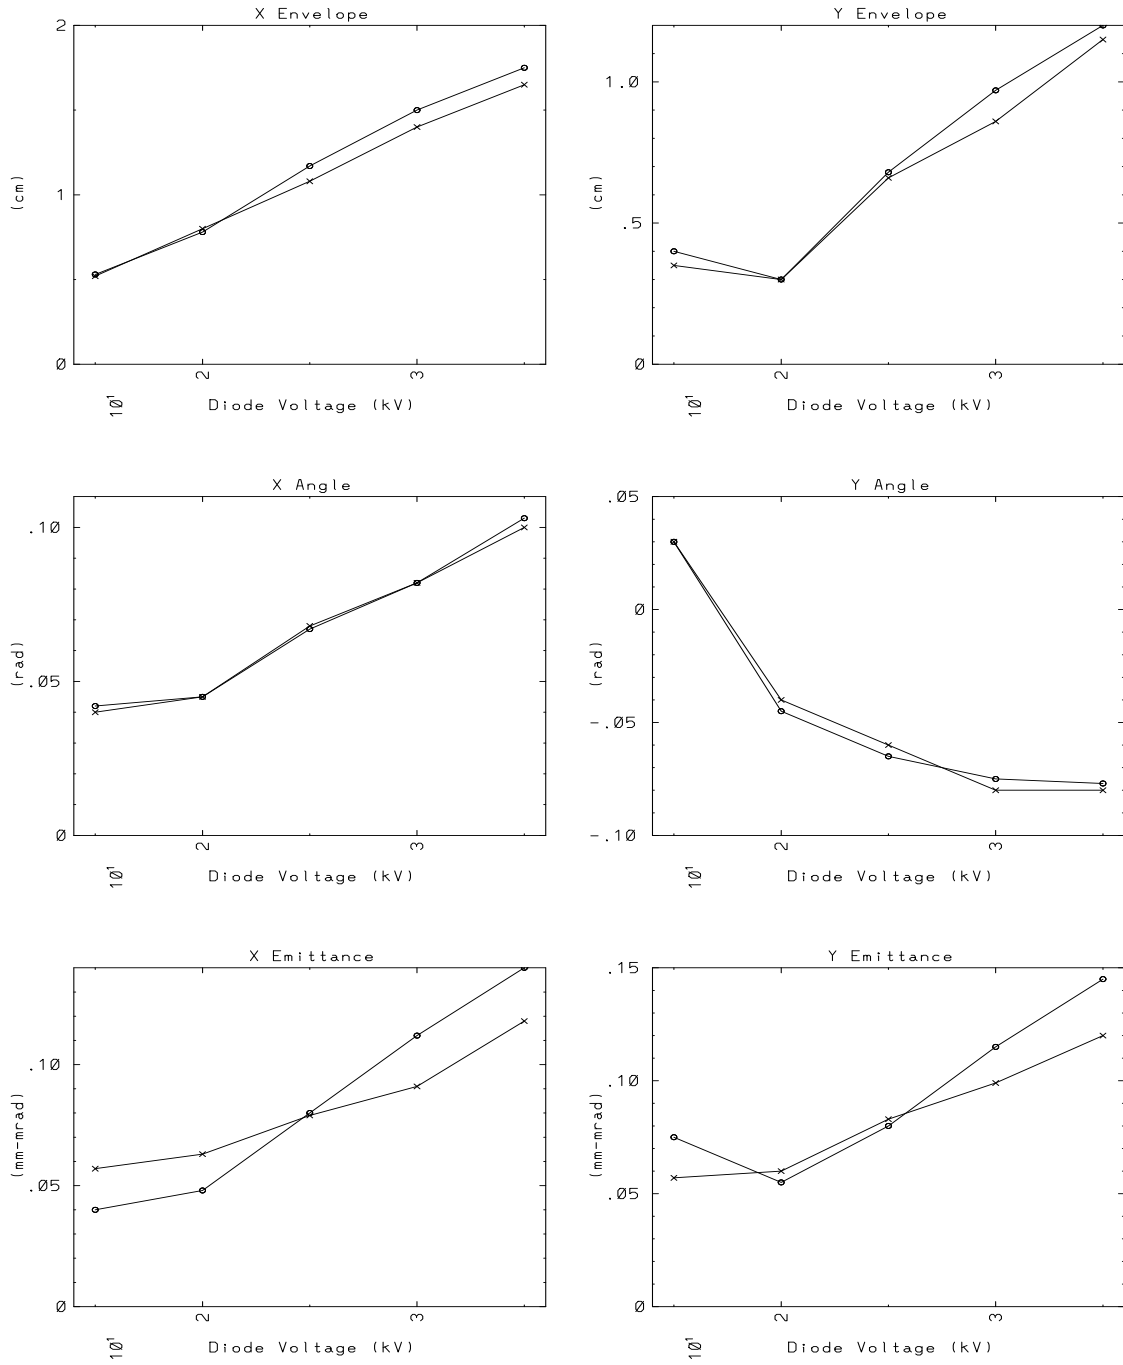


Figure 6.15: Comparison of WARP3d and experiment for the 1 MV scaled injector. The o's are the WARP3d results and the x's are the experimental results.

Y'
(rad)

Y (m)




Figure 6.17: Comparison of experiment and simulation of 30 kV diode injector.

Y'
(rad)

Y (m)




Figure 6.19: Comparison of experiment and simulation of 15 kV diode injector.

Simulations of the scaled injector experiment show good agreement with the experimental results. This shows that the behavior of the beam in the injector is well understood. The agreement also validates the results of WARP3d. The use of WARP3d to design the scaled 1 MV injector experiment demonstrates that the three-dimensional code can serve as a useful design tool.

Chapter 7

Summary

Beam quality is a critical issue for a heavy ion fusion driver because the beam must be focused down onto a small target with a large standoff distance between the final focusing lens and the target. Low transverse and longitudinal temperatures are required. The beam must not be heated too much during the transport from the injector to the target. Making that possible requires understanding of the beam dynamics as the beam is manipulated along the accelerator. The effect of non-linear fields and misalignments of accelerator elements must be understood. All of this must be done including the large space-charge self-fields. Any non-linear or misaligned fields can give a non-linear space-charge distribution which can thermalize to a uniform distribution and a heated beam.

The necessity of including the high space-charge self-fields makes plasma methods, in particular the particle-in-cell method, ideal. Also, the fields from accelerator elements need to be included. Many of the manipulations that the beam goes through are inherently three dimensional; the interaction between transverse and longitudinal must be included. These considerations all led to the development of WARP, a multiple dimensional PIC code containing an accelerator lattice model, and in particular WARP3d, the three-dimensional package of WARP.

The simulations described in Chapters 3 and 4 are mainly test pieces for WARP3d. Chapter 3 shows that there are no serious numerical problems, such as heating or

instabilities, and that the code, as designed and built, is fast enough for our purposes. A three-dimensional, high current ion beam has never been propagated nearly as far with other codes. The main physics obtained from the results of Chapter 3 is the approach of the ends of the beam to equilibrium and the equilibration between transverse and longitudinal thermal energy. Both issues bear further study but are best handled with lower dimensional codes, e.g. WARPrz, especially the evolution of the ends of the beam since it is a slow process. Chapter 4 examines a possible numerical problem arising from the shape of the surrounding conducting pipe as modeled and its image fields (from the beam space-charge). The results show that the effect of the image charges induced on the pipe is insensitive to the pipe shape.

The two applications described in Chapters 5 and 6 elucidate the damaging effects of non-linear external fields, especially when the beam fills a significant fraction of the channel. When a beam extends out into a region of large non-linear fields, the trace space is distorted. If uncorrected, this distortion will thermalize, heating the beam transversely, so that it becomes more difficult to focus.

When a beam is compressed axially, the charge density will remain roughly constant (with constant focusing), forcing the beam to a larger transverse size. The results described in Chapter 5 show that the beam can expand enough to be affected by non-linear fields. This is of concern in the design of a driver for heavy ion fusion where a low emittance is required. The compression must be done gently enough that the focusing can be increased to counteract the transverse expansion.

In an injector, when the beam energy is low, large focusing fields are needed to counteract the beam self-fields. The desired voltages can exceed the limits imposed by voltage breakdown. When that occurs, a larger beam size must be dealt with by minimizing the non-linear fields. The results of Chapter 6 show how the effect of the non-linear fields can be, and indeed are, reduced.

Numerous other applications of WARP3d to heavy ion fusion are anticipated. In addition, the numerical methods developed in the course of this thesis research should prove useful for other applications as well.

Appendix A

Envelope Calculation

The envelope package of the WARP code calculates the envelope of a beam in a lattice of hard edged, uniform quadrupoles (either electric quadrupoles, or magnetic quadrupoles) or in a uniform electric focusing field. It is principally used by the particle loading routines for the transverse particle distribution.

The envelope is calculated in both the x and y directions. The equations for the two directions are coupled by the self-field term. The envelope equations solved are[14]

$$\begin{aligned} a'' &= k_x a + \frac{2K}{a+b} + \epsilon^2/a^3 \\ b'' &= k_y b + \frac{2K}{a+b} + \epsilon^2/b^3 \end{aligned} \tag{A.1}$$

where a and b are the x and y envelopes semi-axis, k_x and k_y are the focusing fields in the appropriate direction, K is the generalized perveance, and ϵ is the unnormalized emittance. The k_x and k_y can be any one of the focusing fields. Note that with either of the quadrupole fields, $k_y = -k_x$, and with uniform focusing $k_y = k_x$. The k_x and k_y are nonzero within a quadrupole, and zero outside. For uniform focusing, k_x and k_y are nonzero everywhere. For electric quadrupole focusing, k_x (and $-k_y$) is given by

$$k_x = -\frac{Ze}{A\alpha\gamma v^2} \frac{dE_x}{dx}, \tag{A.2}$$

where Z is the charge state of the beam ions, e is the charge of an electron, A is the mass of the ion in atomic mass units, α is the conversion factor from atomic mass units to kilograms, v is the beam velocity, γ is the usual relativistic factor, and E_x is the electric focusing field in the x direction. For magnetic quadrupole focusing, k_x (and $-k_y$) is given by

$$k_x = -\frac{Ze}{A\alpha\gamma v} \frac{dB_y}{dx}, \tag{A.3}$$

where B_y is the magnetic focusing field in the y direction. For uniform focusing, k_x (and k_y) is given by

$$k_x = \frac{Ze}{A\alpha\gamma v^2} \frac{dE_r}{dr}, \quad (\text{A.4})$$

where E_r is the radial electric field, and r is the radius. The generalized perveance is given by

$$K = \frac{IZe}{2\pi\epsilon_0 A\alpha(v\gamma)^3} \quad (\text{A.5})$$

where I is the beam current, and ϵ_0 is the permittivity of free space. In equations A.2 through A.5, all of the quantities on the right hand sides, as well as the emittance, are input quantities.

The envelope is calculated starting from initial values. From there, the envelope is advanced with the isochronous leap frog method with z as the independent variable. The steps for the x envelope are

$$\begin{aligned} a'_{i+\frac{1}{2}} &= a'_i + \frac{1}{2}F(a_i)dz \\ a_{i+1} &= a_i + a'_{i+\frac{1}{2}}dz \\ a'_{i+1} &= a'_{i+\frac{1}{2}} + \frac{1}{2}F(a_{i+1})dz \end{aligned} \quad (\text{A.6})$$

where the subscript refers to time level, F is the total force (including focusing, self-fields, and emittance), and dz is the step size in z . The steps for the y envelope are the same but advance b and b' . At the entrance and exit of a quadrupole, the focusing field that is applied to the envelope is adjusted by the amount of the step within the quadrupole. The idea is the same as residence correction done in the particle advance.

Along with the envelope, four test particles are advanced through the focusing fields. Two of the particles do not feel the self-fields. They are used to calculate the undepressed tune shift. Two do feel the self-fields. They are used to calculate the depressed tune shift. For both pairs, one of the particles has the same starting point as the envelope, the other starts on axis with a transverse angle of 1 radian. The transfer matrix calculated from these points is used to calculate the tune shift. If (x_0, x'_0) is the starting point, and (x_f, x'_f) is a final point, the following equation can be written.

$$\begin{pmatrix} x_0 \\ x'_0 \end{pmatrix} = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix} \begin{pmatrix} x_f \\ x'_f \end{pmatrix} \quad (\text{A.7})$$

The matrix M is the transfer matrix. The transfer matrix can be solved for with the above equation known for two particles (two particles are needed to give four equations to calculate the four matrix elements). Once the transfer matrix is known, the tune shift, σ , can be calculated since

$$\cos \sigma = \frac{1}{2}\text{Tr}(M) = \frac{1}{2}(m_{11} + m_{22}), \quad (\text{A.8})$$

where Tr means trace. Derivation of this relation can be found in most books on accelerator physics, for example reference [14].

Generally, it is desirable to load a beam that is matched, i.e. with the envelope parameters the same after one lattice period. As a check on how well-matched the envelope is, the differences of the envelope parameters after one lattice period are calculated and printed. The code offers no method to directly calculate a set of matched parameters, but does have a multidimensional minimization routine that searches for matched parameters. The method used is the “amoeba algorithm” from [11].

Appendix B

Behavior of a Beam in Pure Non-Linear Fields

The behavior of a beam in pure, non-linear fields was examined as an aside to the ESQ injector simulations. The beams in the presence of fourth order fields and the dodecapole fields were simulated.

The fourth order field whose effect is simplest to describe is the V_{42} fields. The potential and electric field are given by

$$\begin{aligned}\phi_{42} &= V_{42}r^4 \cos 2\theta = V_{42}(x^4 - y^4) \\ E_x &= -V_{42}4x^3 \\ E_y &= +V_{42}4y^3.\end{aligned}\tag{B.1}$$

The x electric field is independent of y so particles at all y experience the same deflection in x . Likewise for the y electric field which is independent of x . The effect on trace space is to produce a bent curve. Figure B.1 shows the effect on a beam which is very cold in the transverse direction. Trace space remains a line but particles with positive x are pushed negative, particles with negative x are pushed positive, and vice versa for y .

With the components V_{40} and V_{44} , the fields are not independent of the perpendicular direction as in V_{42} . The pure octopole potential and fields are given by

$$\begin{aligned}\phi_{44} &= V_{44}r^4 \cos 4\theta = V_{44}(x^4 + y^4 - 6x^2y^2) \\ E_x &= V_{44}(12xy^2 - 4x^3) \\ E_y &= V_{44}(12yx^2 - 4y^3).\end{aligned}\tag{B.2}$$

Here, the x electric field is a function of both x and y , so particles with varying y will be in a different field. This has the effect of spreading the distribution in trace space. Figure B.2 shows the effect of a pure octopole field on a cold beam. The particles along the x axis ($y = 0$) feel the same field as from the V_{42} field and so lie on the same

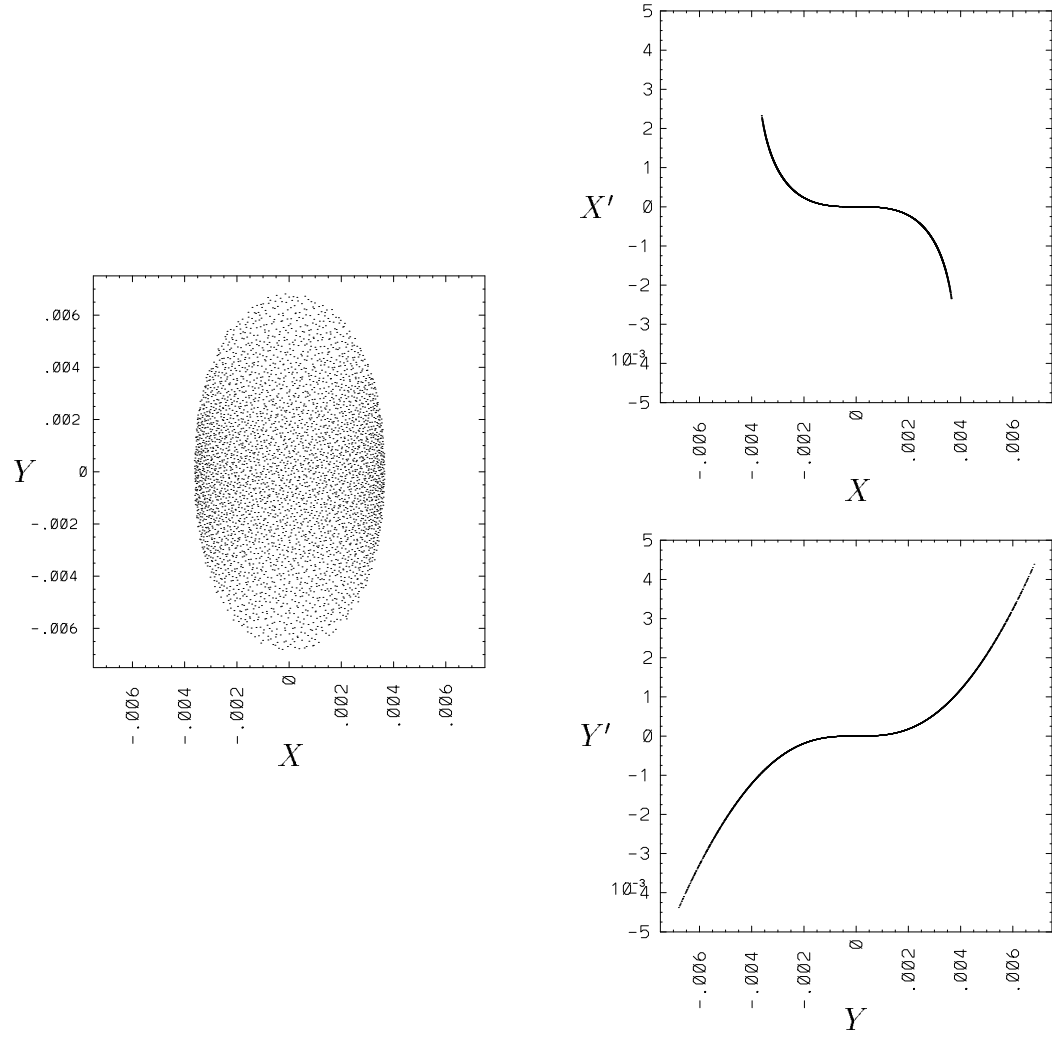


Figure B.1: Effect of V_{42} on a cold beam.

line as in $x - x'$ space in B.1. For particles which have nonzero y , though, the electric field is diminished by the term $12xy^2$ and so are not deflected as much. Further away from the x axis, the term $12xy^2$ becomes greater than $4x^3$, and particles are deflected the other way. The “lobes” cross the x axis.

The potential and fields of V_{40} are given by

$$\begin{aligned}\phi_{40} &= V_{40}r^4 = V_{40}(x^4 + 2x^2y^2 + y^4) \\ E_x &= V_{40}(-4x^3 - 4xy^2) \\ E_y &= V_{40}(-4y^3 - 4yx^2).\end{aligned}\tag{B.3}$$

Again, the x electric field is a function of both x and y , but now, both terms have the same sign. Figure B.3 shows the effect on a cold beam. The particles along the x axis follow the same line as above in $x - x'$ space. Particles off the x axis have the field increased by the $4yx^2$ term and are deflected more. A straight line appears in trace space. These are the particles that are on the edge of the beam. Their position in y is given by $(a^2 - x^2)^{\frac{1}{2}}$, so

$$E_x = V_{40}(-4x^3 - 4x(a^2 - x^2))\tag{B.4}$$

which reduces to $E_x = V_{40}(-4xa^2)$. The field is linear in x .

Finally, the behavior of a beam in a dodecapole field was examined. The potential and fields are

$$\begin{aligned}\phi_{66} &= V_{66}r^6 \cos 6\theta = V_{66}(x^6 - 15x^4y^2 + 15x^2y^4 - y^6) \\ E_x &= V_{66}(-6x^5 + 60x^3y^2 - 30xy^4) \\ E_y &= V_{66}(6y^5 - 60y^3x^2 + 30yx^4).\end{aligned}\tag{B.5}$$

The effect of a pure dodecapole field is shown in figure B.4. The spreading is complicated. The transverse profile forms a nice hexagon.

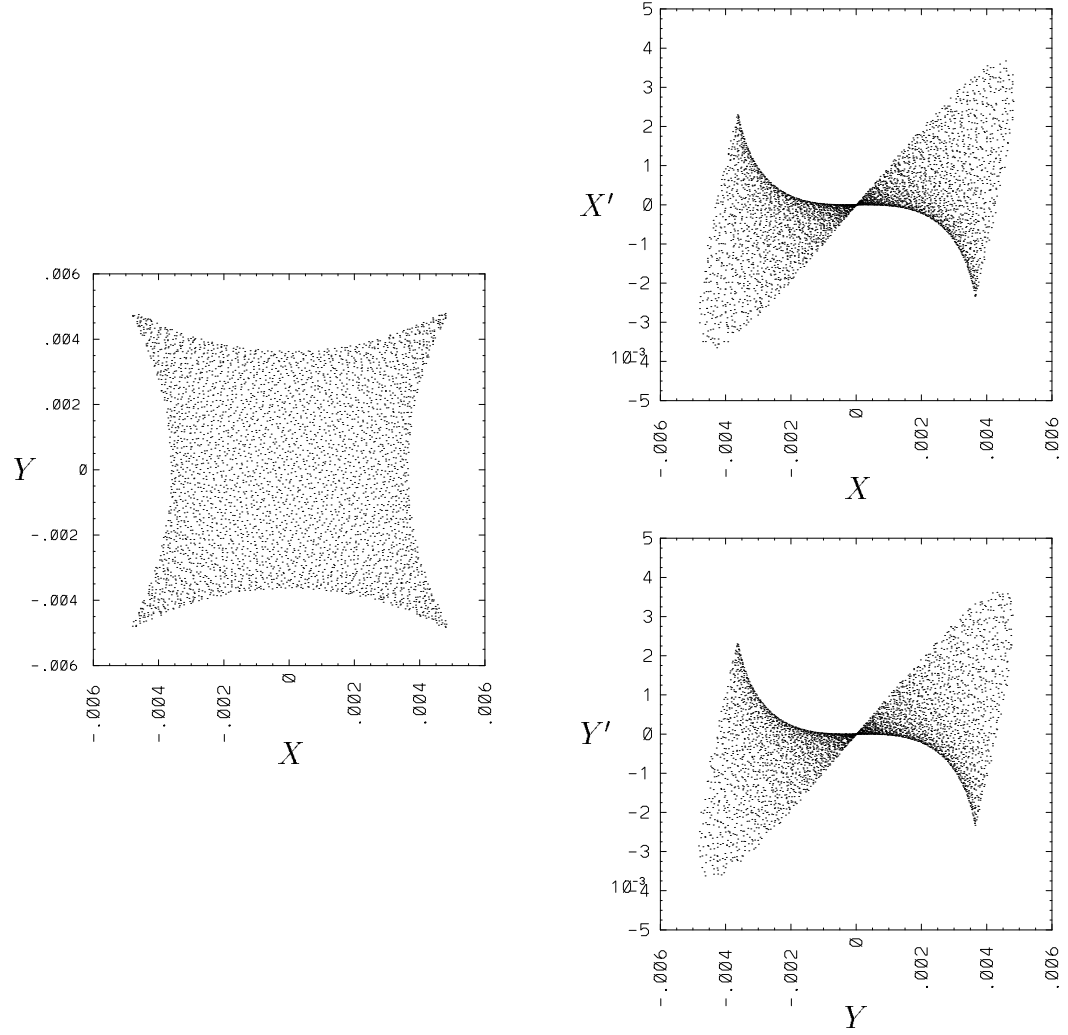


Figure B.2: Effect of V_{44} on a cold beam.




Figure B.3: Effect of V_{40} on a cold beam.

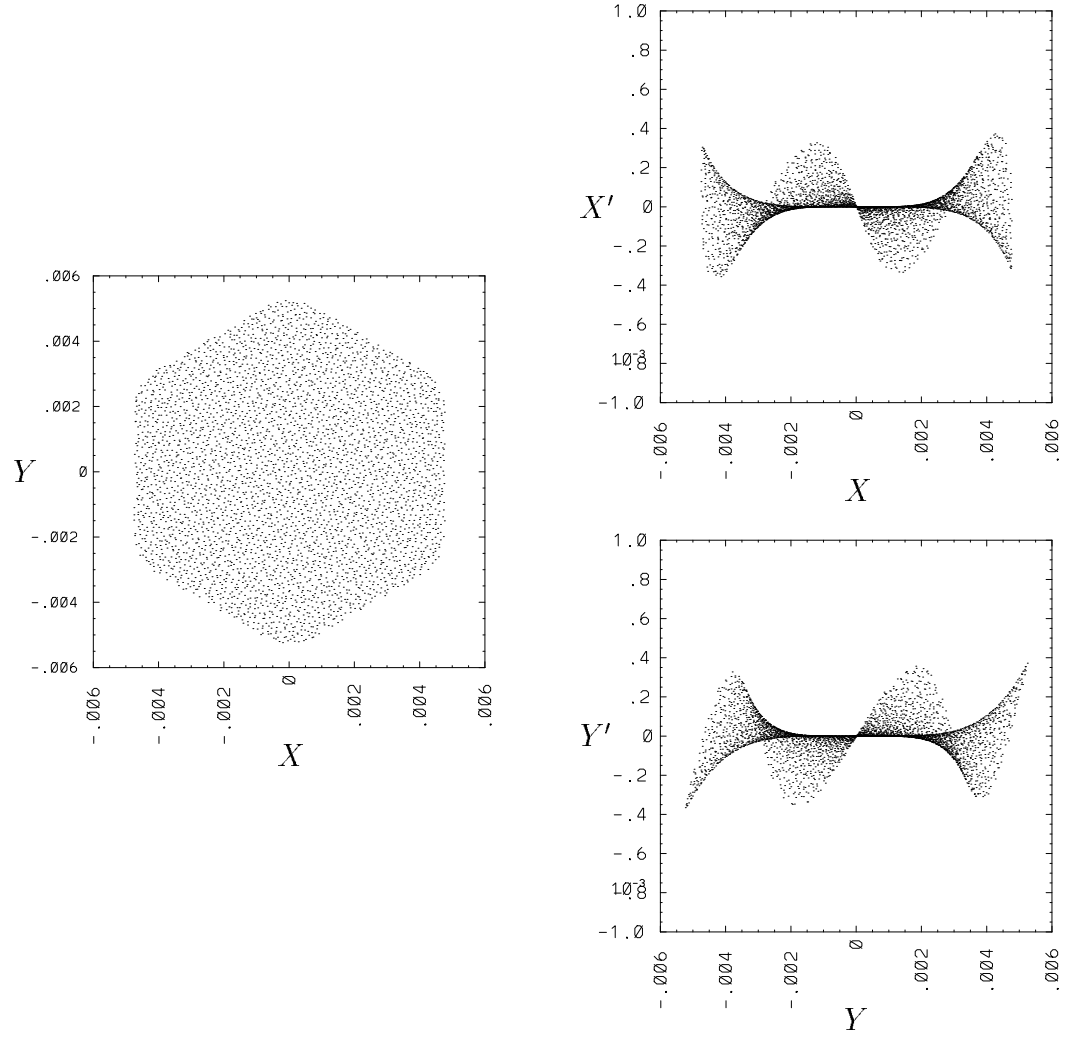


Figure B.4: Effect of V_{66} on a cold beam.

Bibliography

- [1] A. W. Maschke. “Initiation of thermonuclear detonations.” Technical Report BNL 19008, Brookhaven National Laboratory, June 1974.
- [2] R. T. Avery. “Program plan for the MBE-4 multiple beam experiment.” Technical Report 19592, LBL, 1985.
- [3] T. Fessenden, R. Bangerter, D. Berners, J. Chew, S. Eylon, A. Faltens, W. Fawley, C. Fong, M. Fong, K. Hahn, E. Henestroza, D. Judd, E. Lee, C. Lionberger, S. Mukherjee, C. Peters, C. Pike, G. Raymond, L. Reginato, H. Rutkowski, P. Seidl, L. Smith, D. Vanecek, S. Yu, F. Deadrick, A. Friedman, L. Griffith, D. Hewett, M. Newton, and H. Shay. “ILSE, the next step toward a heavy ion induction accelerator for inertial fusion energy.” In *Proceedings of the International Conference on Plasma Physics and Controlled Nuclear Fusion Research*, Wurzburg, Germany, September 1992. IAEA.
- [4] I.M. Kapchinskij and V.V. Vladimirskij. “Limitations of proton beam current in a strong focusing linear accelerator associated with the beam space charge.” In *Proceedings of the Second International Conference on High Energy Accelerators*, pages 274–288, 1959.
- [5] Charles K. Birdsall and A. Bruce Langdon. *Plasma Physics via Computer Simulation*. McGraw-Hill Book Company, 1985.
- [6] R. W. Hockney and J. W. Eastwood. *Computer Simulations Using Particles*. Adam Hilger, Bristol, 1988.
- [7] A. Friedman and S. P. Auerbach. “Numerically induced stochasticity.” *Journal of Computational Physics*, 93:171, 1991.
- [8] S. K. Zaremba. “La méthode des bons treillis pour le calcul des intégrales multiples.” In S. K. Zaremba, editor, *Applications of Number Theory to Numerical Analysis*, pages 39–119. Academic Press, 1972.

- [9] H. Stephen Au-Yeung, Yu-Jiuan Chen, Charles K. Birdsall, and A. Bruce Langdon. "Quiet start method comparisons; random, bit reversed, and Fibonacci numbers." Fourth quarterly progress report, Electronics Research Laboratory, University of California, Berkeley, 1980.
- [10] David Neuffer. "Longitudinal motion in high current ion beams - A self-consistent phase space distribution with an envelope equation." *IEEE Transactions on Nuclear Science*, NS26(3):3031, June 1979.
- [11] William H. Press et. al. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 1986. pp. 652-659.
- [12] Dale E. Nielsen, John J. Ambrosiano, Alex Friedman, and A. Bruce Langdon. "Resistive block methods for electrostatic field solutions with internal elements." In *Proceedings of the Thirteenth Conference on the Numerical Simulation of Plasmas*, 1989. Final two authors added after submission of abstract.
- [13] Paul F. Dubois et. al. "The Basis system." Technical Report M-225, LLNL, 1988.
- [14] Henri Bruck. *Accélérateurs Circulaires de Particules*. Presses Universitaires de France, 1966. English translation available through Las Alamos Laboratory, paper number LA-TR-72-10 Rev.
- [15] D.A. Callahan. "Simulations of longitudinal beam dynamics of space-charge dominated beams for heavy ion fusion." U. C. Davis, 1993. (Doctoral thesis)
- [16] S. Eylon et. al. "Transverse beam dynamics studies with the mbe-4 multiple beam, heavy-ion linear accelerator." *Journal of Applied Physics*. (in preparation)
- [17] Shmuel Eylon et. al. "Drift compression experiments on mbe-4 and related emittance growth phenomena." In *Proceedings of the Fourteenth Biennial IEEE/APS Particle Accelerator Conference*, May 1991.
- [18] C. M. Celata, I. Haber, L. Smith, and M. G. Teifenback. "The effect of induced charge at boundaries on the transverse dynamics of a space-charge-dominated beam." *IEEE Transactions on Nuclear Science*, NS-32:2480, 1985.
- [19] Alex Friedman, David P. Grote and Irving Haber. "Three-dimensional particle simulation of heavy-ion fusion beams." *Phys. Fluids B*, July 1992, 4-7. pp. 2203-2210.
- [20] Alex Friedman, Debra A. Callahan, David P. Grote, A. Bruce Langdon and Irving Haber. "WARP: A 3D (+) PIC code for HIF simulations" In *Proceedings*

- of Conference on Computer Codes and the Linear Accelerator Community*, LANL report LA-11857-C, January 1990. pp. 21–25, Proceedings edit by R.K. Cooper and K. C. D. Chan.
- [21] David P. Grote et. al. “3D Simulations of Axially Confined Heavy Ion Beams in Round and Square Pipes” In *Proceedings of the International Symposium on Heavy Ion Inertial Fusion*, December, 1990.
- [22] Alex Friedman, David P. Grote, Debra A. Callahan, A. Bruce Langdon and Irving Haber. “3d Particle Simulation of Beams Using the WARP Code: Transport Around Bends” *Particle Accelerators*, 1992, 37–38. page 131.
- [23] David P. Grote, Alex Friedman and Irving Haber”, “3d Simulations of Axially Confined Heavy Ion Beams in Round and Square Pipes” *Particle Accelerators*, 1992, 37–39. page 141.